



US009459865B2

(12) **United States Patent**
Ould-Ahmed-Vall et al.

(10) **Patent No.:** **US 9,459,865 B2**
(45) **Date of Patent:** **Oct. 4, 2016**

(54) **SYSTEMS, APPARATUSES, AND METHODS FOR PERFORMING A BUTTERFLY HORIZONTAL AND CROSS ADD OR SUBTRACT IN RESPONSE TO A SINGLE INSTRUCTION**

9/30036 (2013.01); **G06F 9/30145** (2013.01);
G06F 17/142 (2013.01)

(58) **Field of Classification Search**

CPC **G06F 9/30036**; **G06F 17/142**
USPC **712/222**; **708/404**
See application file for complete search history.

(75) Inventors: **Elmoustapha Ould-Ahmed-Vall**,
Chandler, AZ (US); **Mostafa Hagog**,
Kaukab (IL); **Robert Valentine**, Kiryat
Tivon (IL); **Amit Gradstein**, Binyamina
(IL); **Simon Rubanovich**, Haifa (IL);
Zeev Sperber, Zichron Yackov (IL)

(56)

References Cited

U.S. PATENT DOCUMENTS

5,204,962 A * 4/1993 Uramoto **G06F 17/10**
708/409
5,297,070 A * 3/1994 Hua **G06F 5/015**
708/404
6,088,782 A 7/2000 Lee et al.

(Continued)

(*) Notice: Subject to any disclaimer, the term of this
patent is extended or adjusted under 35
U.S.C. 154(b) by 546 days.

OTHER PUBLICATIONS

Notification of Transmittal of the International Search Report and
the Written Opinion or the International Searching Authority, or the
Declaration for PCT Counterpart Application No. PCT/US2011/
067183, 9 pgs., (Aug. 29, 2012).

(Continued)

(21) Appl. No.: **13/992,236**

(22) PCT Filed: **Dec. 23, 2011**

(86) PCT No.: **PCT/US2011/067183**

§ 371 (c)(1),

(2), (4) Date: **Jun. 6, 2013**

Primary Examiner — Andrew Caldwell

Assistant Examiner — Michael Metzger

(87) PCT Pub. No.: **WO2013/095631**

PCT Pub. Date: **Jun. 27, 2013**

(74) *Attorney, Agent, or Firm* — Nicholson De Vos
Webster & Elliott LLP

(65) **Prior Publication Data**

US 2014/0201502 A1 Jul. 17, 2014

(51) **Int. Cl.**

G06F 9/30 (2006.01)

G06F 17/14 (2006.01)

(52) **U.S. Cl.**

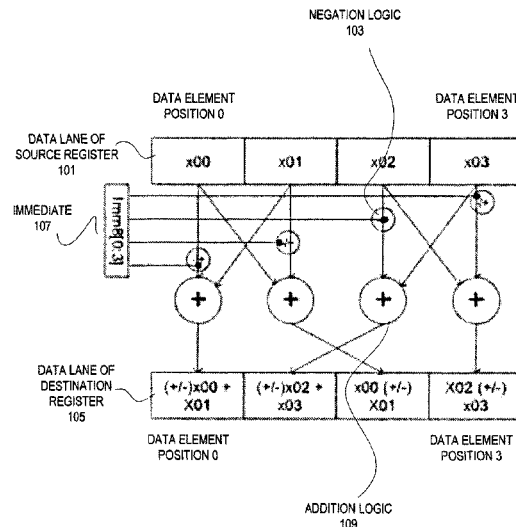
CPC **G06F 9/3001** (2013.01); **G06F 9/30014**
(2013.01); **G06F 9/30018** (2013.01); **G06F**

(57)

ABSTRACT

Embodiments of systems, apparatuses, and methods for
performing in a computer processor vector packed butterfly
horizontal cross add or subtract of packed data elements in
response to a single vector packed butterfly horizontal cross
add or subtract instruction that includes a destination vector
register operand, a source vector register operand, an imme-
diate, and an opcode are described.

19 Claims, 19 Drawing Sheets



(56)

References Cited

OTHER PUBLICATIONS

U.S. PATENT DOCUMENTS

7,095,808	B1	8/2006	Shohara	
2006/0015702	A1 *	1/2006	Khan	G06F 9/30014 712/22
2006/0149938	A1	7/2006	Jiang et al.	
2009/0019269	A1	1/2009	Wolff et al.	

PCT Notification concerning Transmittal of International Preliminary Report on Patentability (Chapter I of the Patent Cooperation Treaty) for PCT Counterpart Application No. PCT/US2011/067183 , 5 pgs., (Jun. 24, 2014).

* cited by examiner

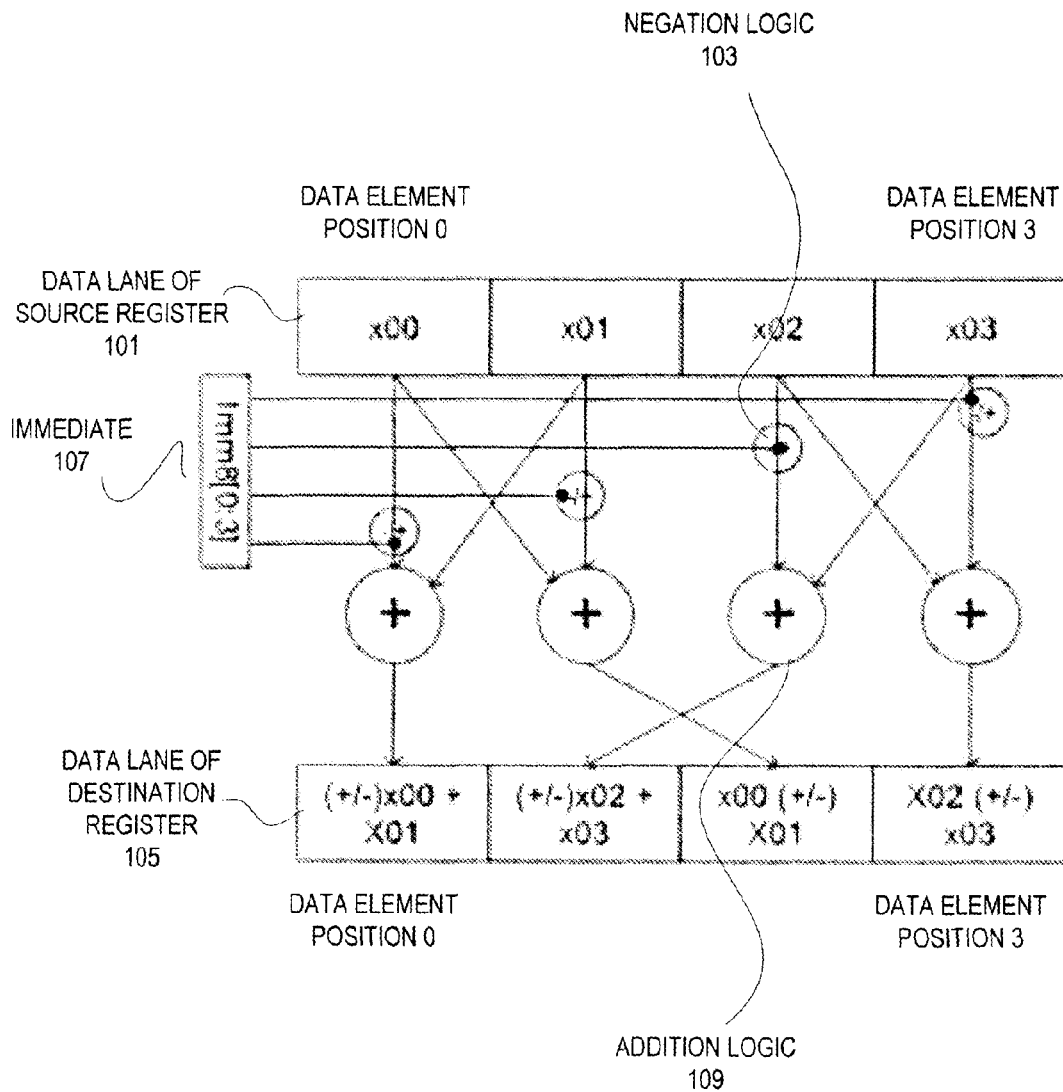


FIGURE 1

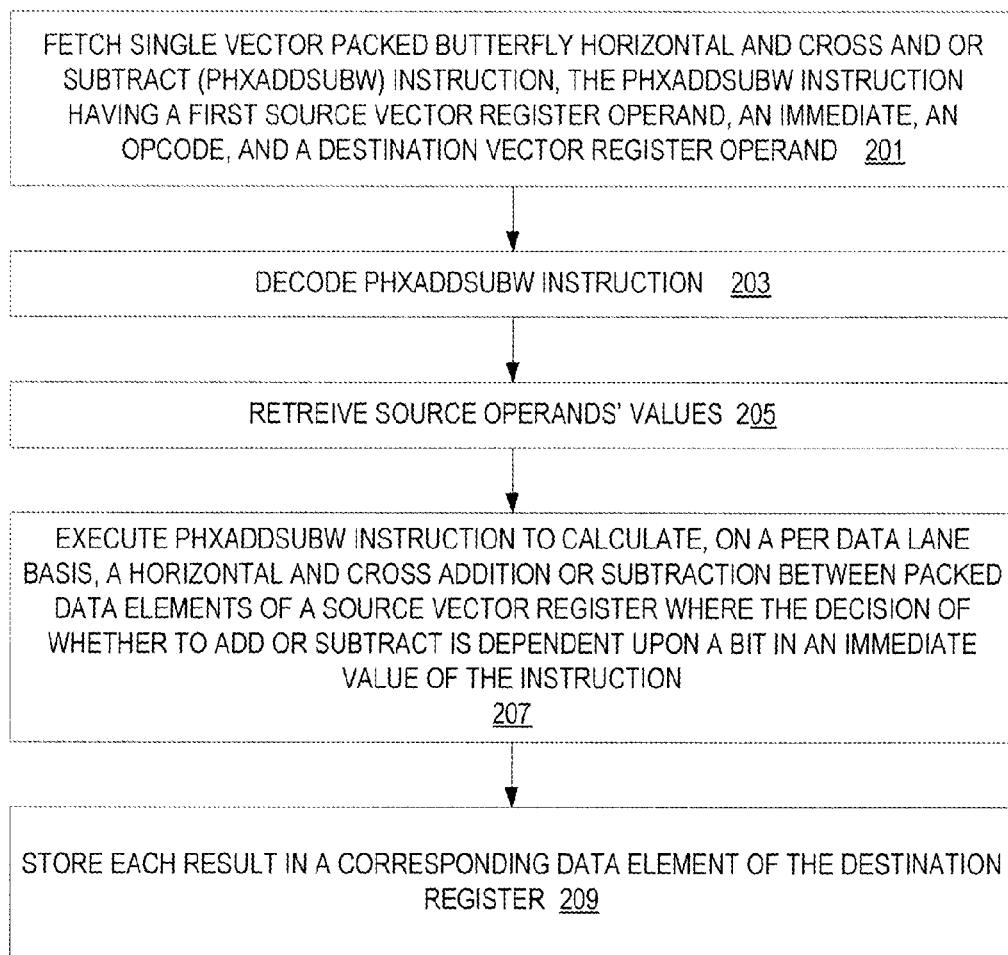


FIGURE 2

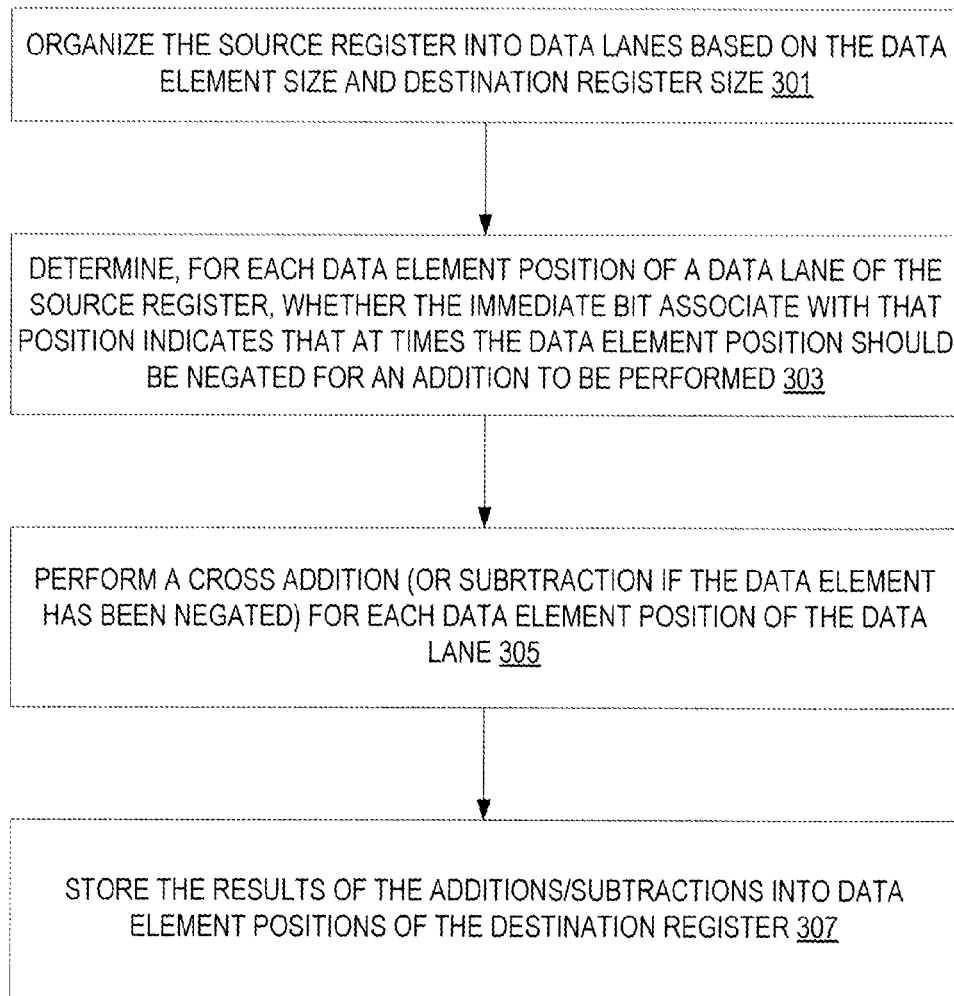


FIGURE 3

Per Data Lane,

$$\text{DEST}[0] = \text{SRC}[0] * (\text{imm}[0] ? -1 : 1) + \text{SRC}[1]$$

$$\text{DEST}[1] = \text{SRC}[2] * (\text{imm}[2] ? -1 : 1) + \text{SRC}[3]$$

$$\text{DEST}[2] = \text{SRC}[1] * (\text{imm}[1] ? -1 : 1) + \text{SRC}[0]$$

$$\text{DEST}[3] = \text{SRC}[3] * (\text{imm}[3] ? -1 : 1) + \text{SRC}[2]$$

FIGURE 4





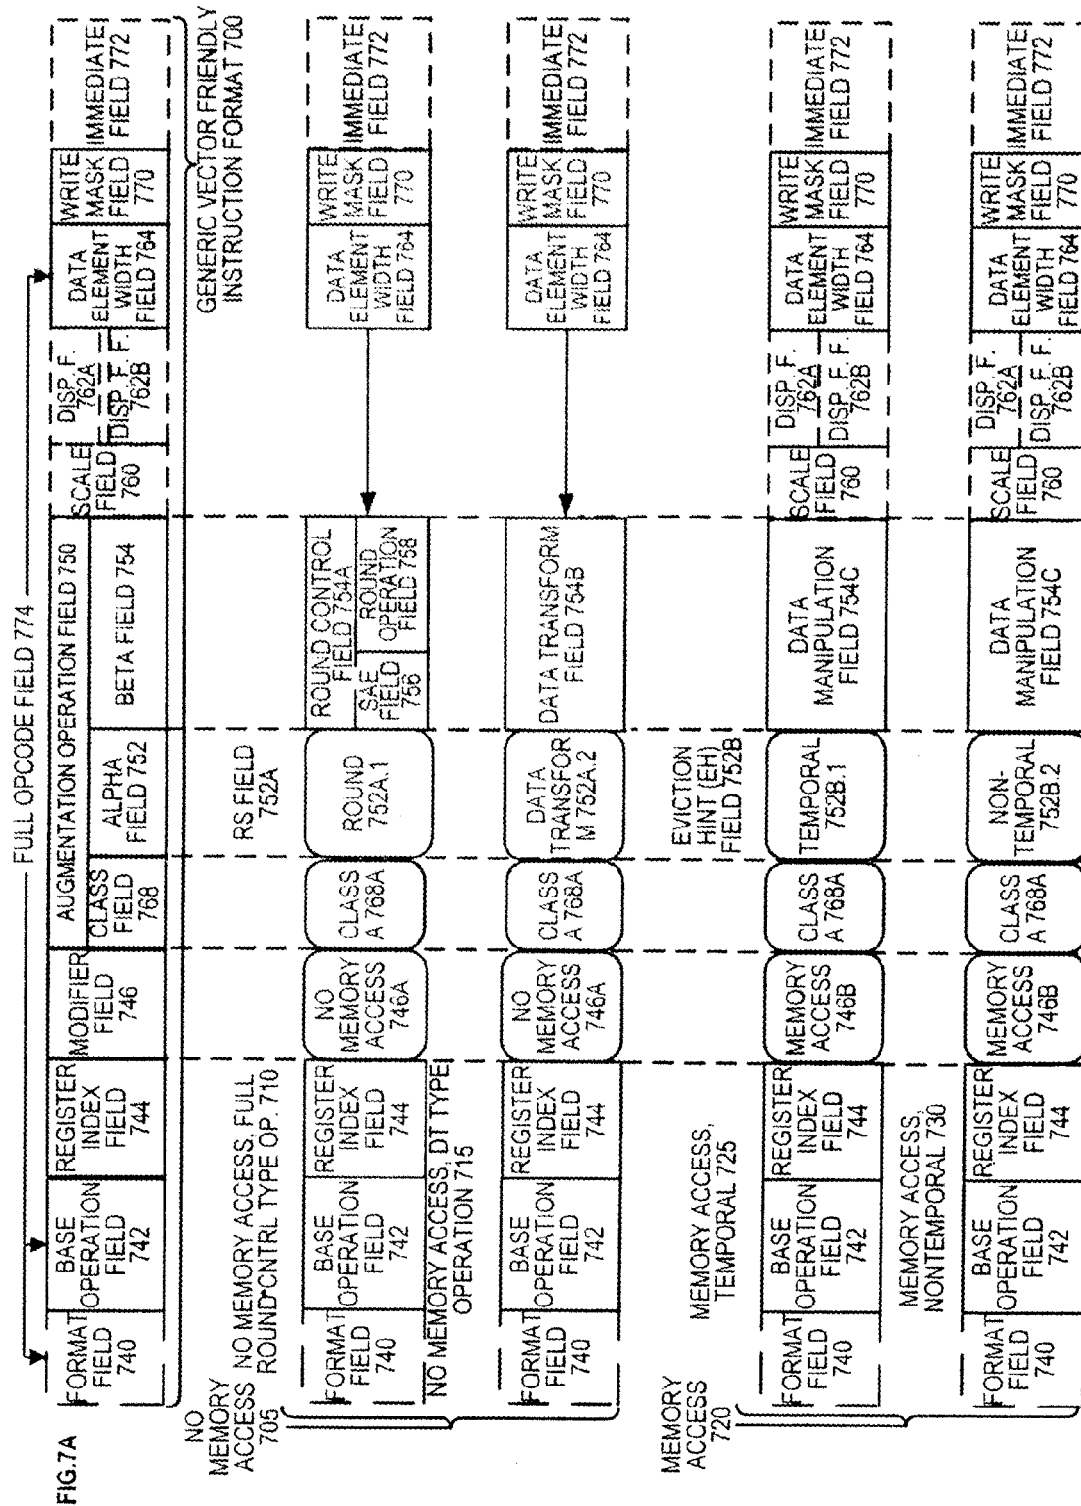
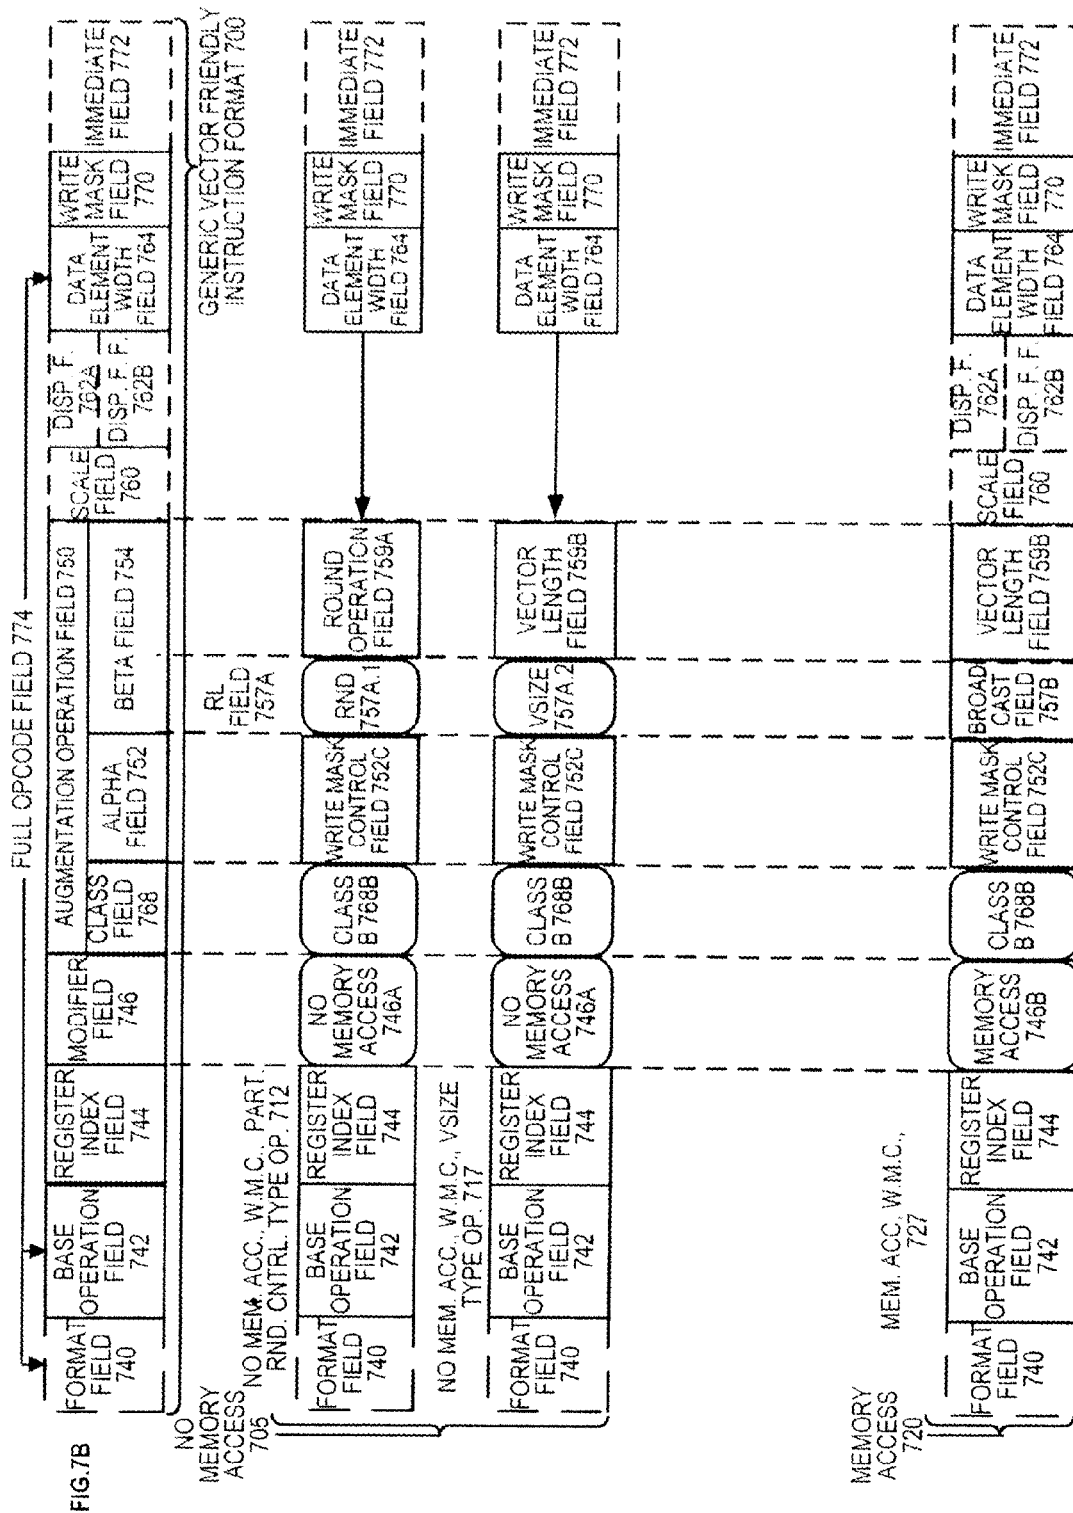





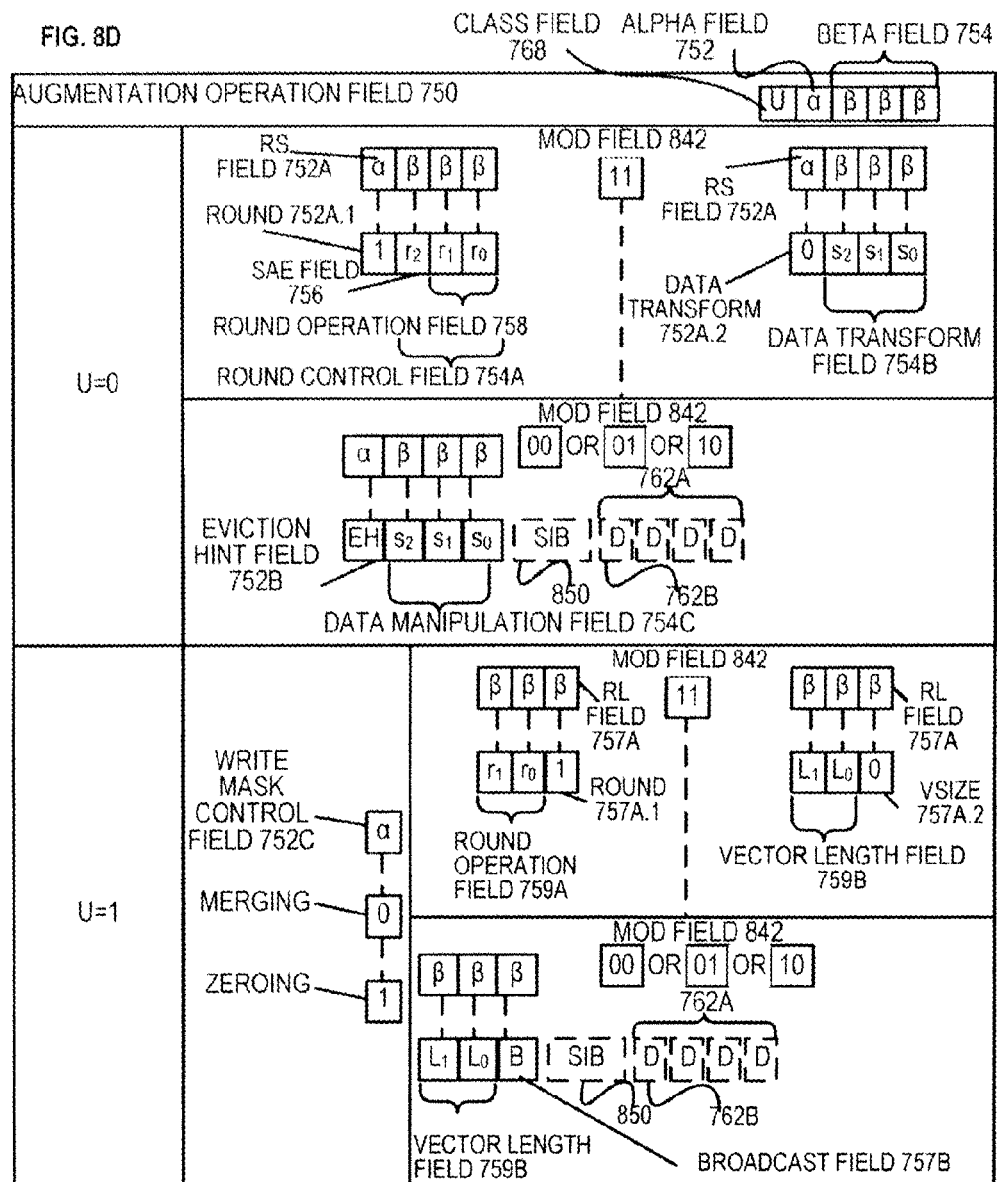
Figure 5




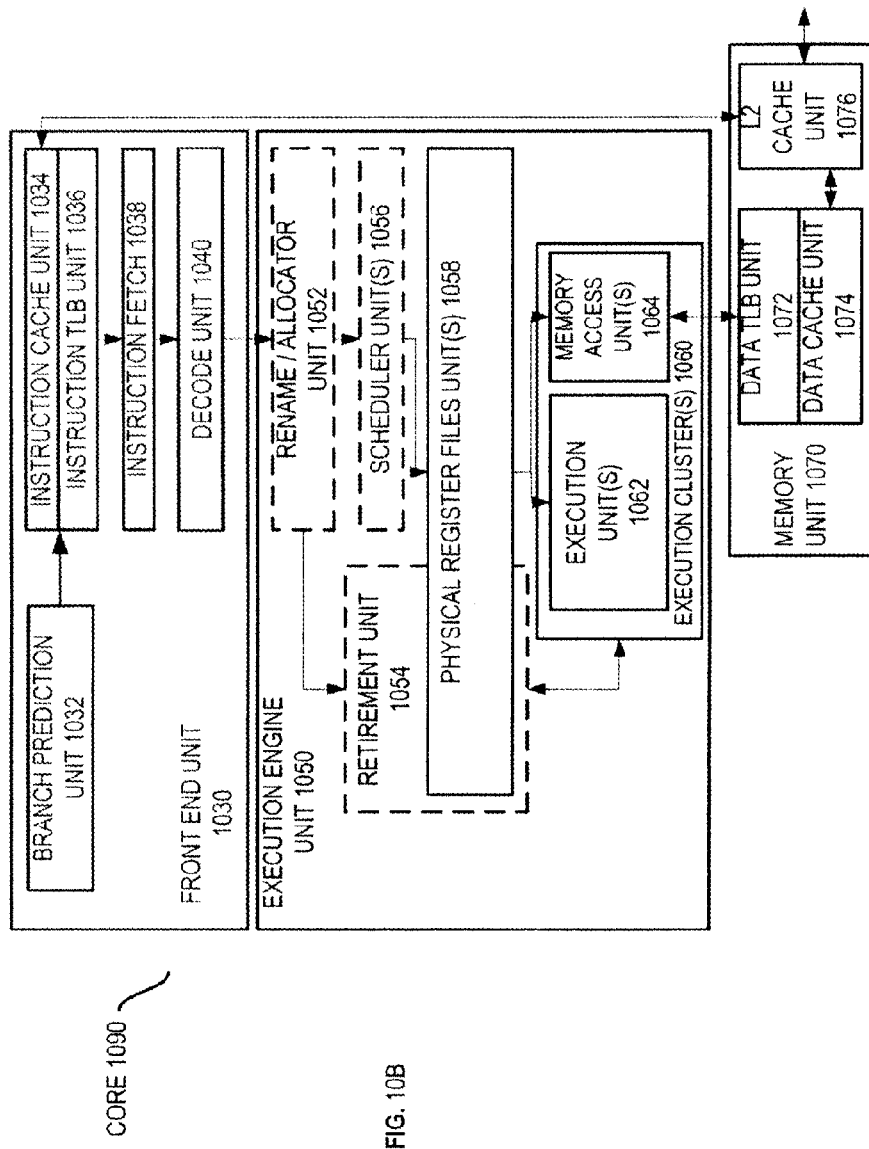



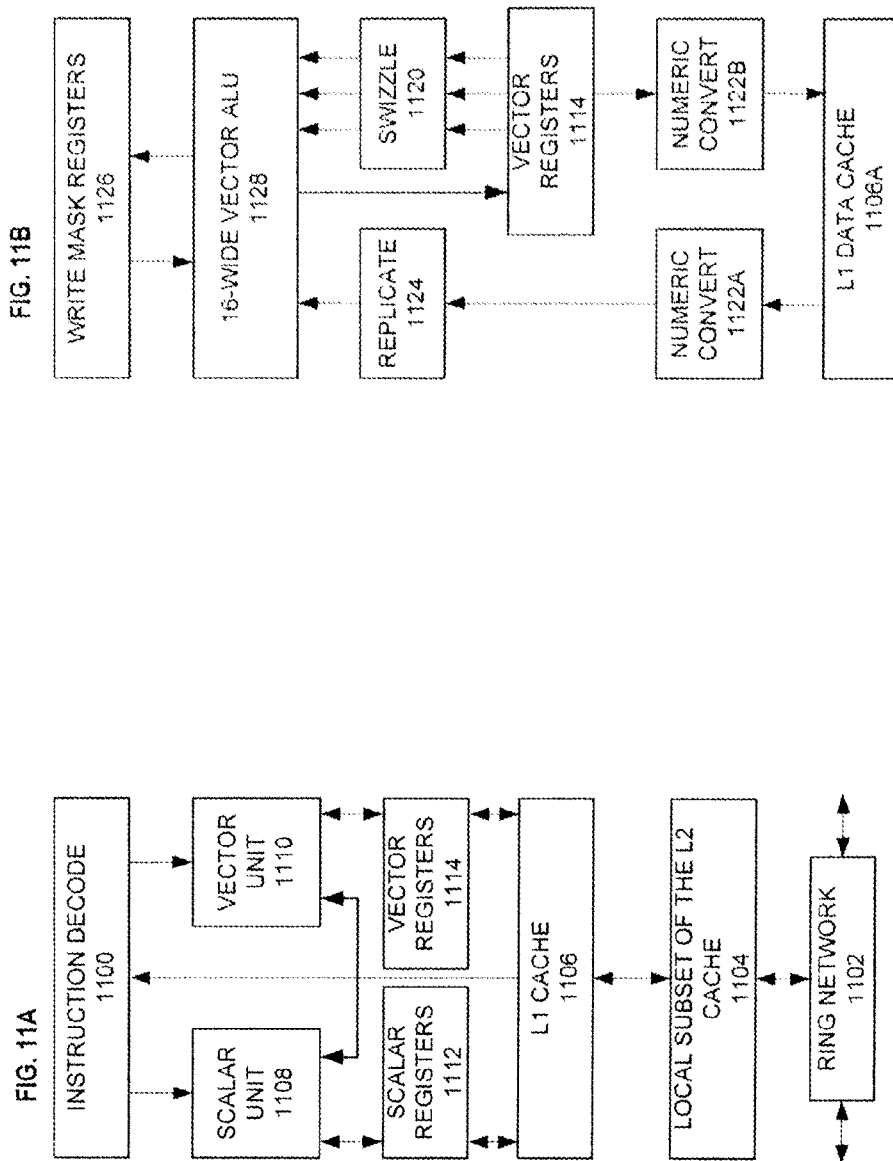


















FIG. 13




FIG. 14




FIG. 15


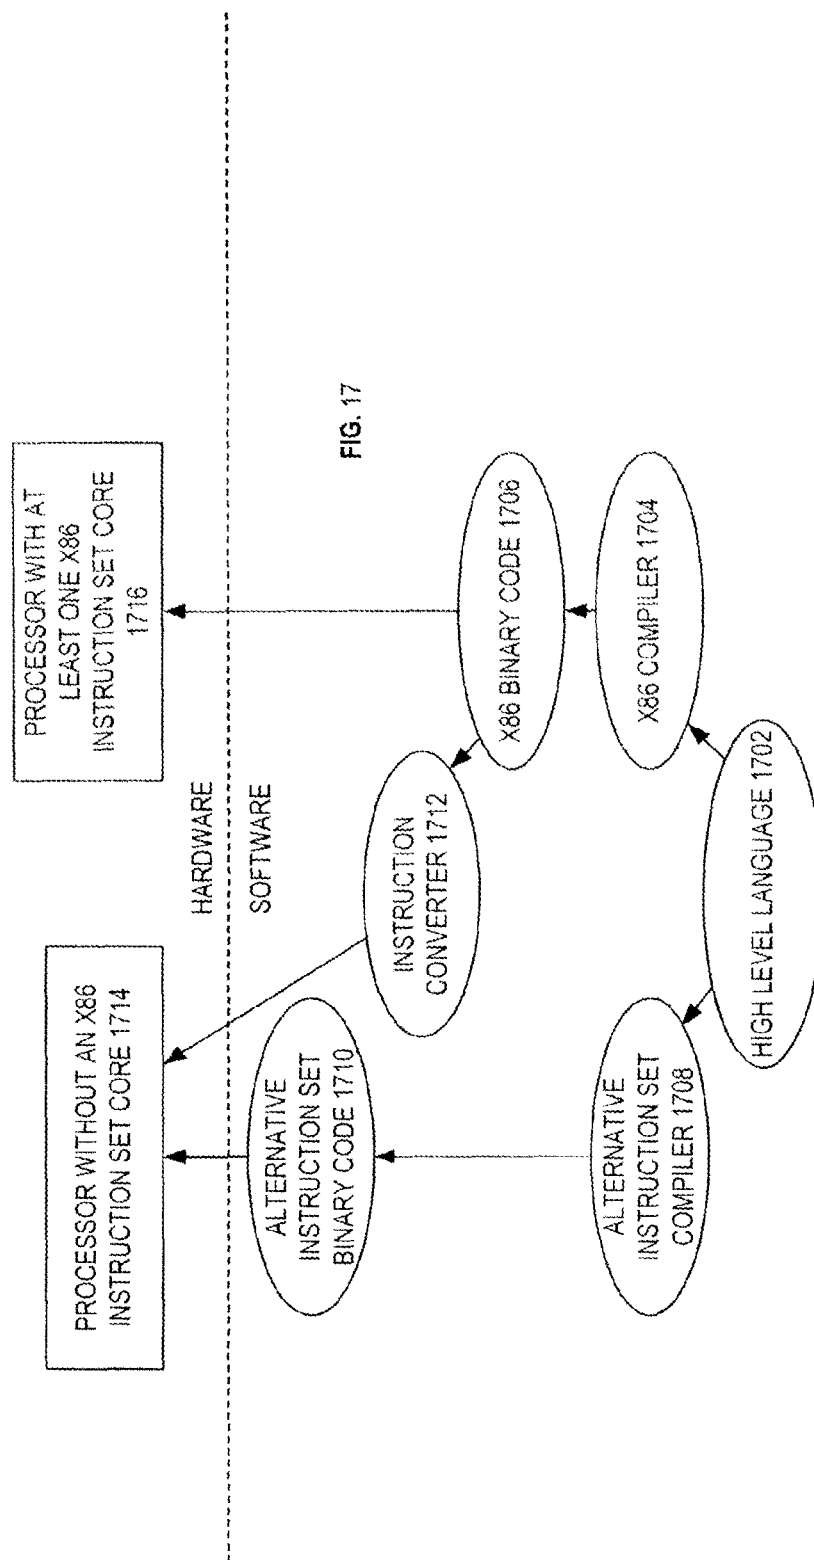


FIG. 16



1

SYSTEMS, APPARATUSES, AND METHODS FOR PERFORMING A BUTTERFLY HORIZONTAL AND CROSS ADD OR SUBTRACT IN RESPONSE TO A SINGLE INSTRUCTION

CROSS-REFERENCE TO RELATED APPLICATIONS

This patent application is a U.S. National Phase Application under 35 U.S.C. §371 of International Application No. PCT/US2011/067183, filed Dec. 23, 2011, entitled SYSTEMS, APPARATUSES, AND METHODS FOR PERFORMING A BUTTERFLY HORIZONTAL AND CROSS ADD OR SUBTRACT IN RESPONSE TO A SINGLE INSTRUCTION.

FIELD OF INVENTION

The field of invention relates generally to computer processor architecture, and, more specifically, to instructions which when executed cause a particular result.

BACKGROUND

An instruction set, or instruction set architecture (ISA), is the part of the computer architecture related to programming, and may include the native data types, instructions, register architecture, addressing modes, memory architecture, interrupt and exception handling, and external input and output (I/O). The term instruction generally refers herein to macro-instructions—that is instructions that are provided to the processor (or instruction converter that translates (e.g., using static binary translation, dynamic binary translation including dynamic compilation), morphs, emulates, or otherwise converts an instruction to one or more other instructions to be processed by the processor) for execution—as opposed to micro-instructions or micro-operations (micro-ops)—that is the result of a processor's decoder decoding macro-instructions.

The ISA is distinguished from the microarchitecture, which is the internal design of the processor implementing the instruction set. Processors with different microarchitectures can share a common instruction set. For example, Intel® Pentium 4 processors, Intel® Core™ processors, and processors from Advanced Micro Devices, Inc. of Sunnyvale Calif. implement nearly identical versions of the x86 instruction set (with some extensions that have been added with newer versions), but have different internal designs. For example, the same register architecture of the ISA may be implemented in different ways in different microarchitectures using well-known techniques, including dedicated physical registers, one or more dynamically allocated physical registers using a register renaming mechanism (e.g., the use of a Register Alias Table (RAT), a Reorder Buffer (ROB), and a retirement register file; the use of multiple maps and a pool of registers), etc. Unless otherwise specified, the phrases register architecture, register file, and register are used herein to refer to that which is visible to the software/programmer and the manner in which instructions specify registers. Where a specificity is desired, the adjective logical, architectural, or software visible will be used to indicate registers/files in the register architecture, while different adjectives will be used to designation registers in a given microarchitecture (e.g., physical register, reorder buffer, retirement register, register pool).

2

An instruction set includes one or more instruction formats. A given instruction format defines various fields (number of bits, location of bits) to specify, among other things, the operation to be performed (opcode) and the operand(s) on which that operation is to be performed. Some instruction formats are further broken down though the definition of instruction templates (or subformats). For example, the instruction templates of a given instruction format may be defined to have different subsets of the instruction format's fields (the included fields are typically in the same order, but at least some have different bit positions because there are less fields included) and/or defined to have a given field interpreted differently. Thus, each instruction of an ISA is expressed using a given instruction format (and, if defined, in a given one of the instruction templates of that instruction format) and includes fields for specifying the operation and the operands. For example, an exemplary ADD instruction has a specific opcode and an instruction format that includes an opcode field to specify that opcode and operand fields to select operands (source1/destination and source2); and an occurrence of this ADD instruction in an instruction stream will have specific contents in the operand fields that select specific operands.

Scientific, financial, auto-vectorized general purpose, RMS (recognition, mining, and synthesis), and visual and multimedia applications (e.g., 2D/3D graphics, image processing, video compression/decompression, voice recognition algorithms and audio manipulation) often require the same operation to be performed on a large number of data items (referred to as “data parallelism”). Single Instruction Multiple Data (SIMD) refers to a type of instruction that causes a processor to perform an operation on multiple data items. SIMD technology is especially suited to processors that can logically divide the bits in a register into a number of fixed-sized data elements, each of which represents a separate value. For example, the bits in a 256-bit register may be specified as a source operand to be operated on as four separate 64-bit packed data elements (quad-word (Q) size data elements), eight separate 32-bit packed data elements (double word (D) size data elements), sixteen separate 16-bit packed data elements (word (W) size data elements), or thirty-two separate 8-bit data elements (byte (B) size data elements). This type of data is referred to as packed data type or vector data type, and operands of this data type are referred to as packed data operands or vector operands. In other words, a packed data item or vector refers to a sequence of packed data elements, and a packed data operand or a vector operand is a source or destination operand of a SIMD instruction (also known as a packed data instruction or a vector instruction).

By way of example, one type of SIMD instruction specifies a single vector operation to be performed on two source vector operands in a vertical fashion to generate a destination vector operand (also referred to as a result vector operand) of the same size, with the same number of data elements, and in the same data element order. The data elements in the source vector operands are referred to as source data elements, while the data elements in the destination vector operand are referred to a destination or result data elements. These source vector operands are of the same size and contain data elements of the same width, and thus they contain the same number of data elements. The source data elements in the same bit positions in the two source vector operands form pairs of data elements (also referred to as corresponding data elements; that is, the data element in data element position 0 of each source operand correspond,

the data element in data element position 1 of each source operand correspond, and so on). The operation specified by that SIMD instruction is performed separately on each of these pairs of source data elements to generate a matching number of result data elements, and thus each pair of source data elements has a corresponding result data element. Since the operation is vertical and since the result vector operand is the same size, has the same number of data elements, and the result data elements are stored in the same data element order as the source vector operands, the result data elements are in the same bit positions of the result vector operand as their corresponding pair of source data elements in the source vector operands. In addition to this exemplary type of SIMD instruction, there are a variety of other types of SIMD instructions (e.g., that has only one or has more than two source vector operands, that operate in a horizontal fashion, that generates a result vector operand that is of a different size, that has a different size data elements, and/or that has a different data element order). It should be understood that the term destination vector operand (or destination operand) is defined as the direct result of performing the operation specified by an instruction, including the storage of that destination operand at a location (be it a register or at a memory address specified by that instruction) so that it may be accessed as a source operand by another instruction (by specification of that same location by the another instruction).

The SIMD technology, such as that employed by the Intel® Core™ processors having an instruction set including x86, MMX™, Streaming SIMD Extensions (SSE), SSE2, SSE3, SSE4.1, and SSE4.2 instructions, has enabled a significant improvement in application performance. An additional set of SIMD extensions, referred to the Advanced Vector Extensions (AVX) (AVX1 and AVX2) and using the Vector Extensions (VEX) coding scheme, has been, has been released and/or published (e.g., see Intel® 64 and IA-32 Architectures Software Developers Manual, October 2011; and see Intel® Advanced Vector Extensions Programming Reference, June 2011).

BRIEF DESCRIPTION OF THE DRAWINGS

The present invention is illustrated by way of example and not limitation in the figures of the accompanying drawings, in which like references indicate similar elements and in which:

FIG. 1 illustrates an exemplary illustration of an operation of an exemplary PHXADDSUB instruction.

FIG. 2 illustrates an embodiment of the use of a PHXADDSUB instruction in a processor.

FIG. 3 illustrates an embodiment of a method for processing a PHXADDSUB instruction.

FIG. 4 illustrates exemplary pseudo-code for a four packed data element destination horizontal add or subtract.

FIG. 5 illustrates a correlation between the number of one active bit vector writemask elements and the vector size and the data element size according to one embodiment of the invention.

FIG. 6A illustrates an exemplary AVX instruction format.

FIG. 6B illustrates which fields from FIG. 6A make up a full opcode field and a base operation field.

FIG. 6C illustrates which fields from FIG. 6A make up a register index field.

FIGS. 7A-7B are block diagrams illustrating a generic vector friendly instruction format and instruction templates thereof according to embodiments of the invention.

FIGS. 8A-D are block diagrams illustrating an exemplary specific vector friendly instruction format according to embodiments of the invention.

FIG. 9 is a block diagram of a register architecture according to one embodiment of the invention.

FIG. 10A is a block diagram illustrating both an exemplary in-order pipeline and an exemplary register renaming, out-of-order issue/execution pipeline according to embodiments of the invention.

FIG. 10B is a block diagram illustrating both an exemplary embodiment of an in-order architecture core and an exemplary register renaming, out-of-order issue/execution architecture core to be included in a processor according to embodiments of the invention.

FIGS. 11A-B illustrate a block diagram of a more specific exemplary in-order core architecture, which core would be one of several logic blocks (including other cores of the same type and/or different types) in a chip.

FIG. 12 is a block diagram of a processor that may have more than one core, may have an integrated memory controller, and may have integrated graphics according to embodiments of the invention.

FIG. 13 is a block diagram of an exemplary system in accordance with an embodiment of the present invention.

FIG. 14 is a block diagram of a first more specific exemplary system in accordance with an embodiment of the present invention.

FIG. 15 is a block diagram of a second more specific exemplary system in accordance with an embodiment of the present invention.

FIG. 16 is a block diagram of a SoC in accordance with an embodiment of the present invention.

FIG. 17 is a block diagram contrasting the use of a software instruction converter to convert binary instructions in a source instruction set to binary instructions in a target instruction set according to embodiments of the invention.

DETAILED DESCRIPTION

In the following description, numerous specific details are set forth. However, it is understood that embodiments of the invention may be practiced without these specific details. In other instances, well-known circuits, structures and techniques have not been shown in detail in order not to obscure the understanding of this description.

References in the specification to “one embodiment,” “an embodiment,” “an example embodiment,” etc., indicate that the embodiment described may include a particular feature, structure, or characteristic, but every embodiment may not necessarily include the particular feature, structure, or characteristic. Moreover, such phrases are not necessarily referring to the same embodiment. Further, when a particular feature, structure, or characteristic is described in connection with an embodiment, it is submitted that it is within the knowledge of one skilled in the art to affect such feature, structure, or characteristic in connection with other embodiments whether or not explicitly described.

Overview

In the description below, there are some items that may need explanation prior to describing the operations of this particular instruction in the instruction set architecture. One such item is called a “writemask register” which is generally used to predicate an operand to conditionally control per-element computational operation (below, the term mask register may also be used and it refers to a writemask register such as the “k” registers discussed below). As used below, a writemask register stores a plurality of bits (16, 32, 64,

etc.) wherein each active bit of the writemask register governs the operation/update of a packed data element of a vector register during SIMD processing. Typically, there is more than one writemask register available for use by a processor core.

The instruction set architecture includes at least some SIMD instructions that specify vector operations and that have fields to select source registers and/or destination registers from these vector registers (an exemplary SIMD instruction may specify a vector operation to be performed on the contents of one or more of the vector registers, and the result of that vector operation to be stored in one of the vector registers). Different embodiments of the invention may have different sized vector registers and support more/less/different sized data elements.

The size of the multi-bit data elements specified by a SIMD instruction (e.g., byte, word, double word, quad word) determines the bit locations of the “data element positions” within a vector register, and the size of the vector operand determines the number of data elements. A packed data element refers to the data stored in a particular position. In other words, depending on the size of the data elements in the destination operand and the size of the destination operand (the total number of bits in the destination operand) (or put another way, depending on the size of the destination operand and the number of data elements within the destination operand), the bit locations of the multi-bit data element positions within the resulting vector operand change (e.g., if the destination for the resulting vector operand is a vector register, then the bit locations of the multi-bit data element positions within the destination vector register change). For example, the bit locations of the multi-bit data elements are different between a vector operation that operates on 32-bit data elements (data element position 0 occupies bit locations 31:0, data element position 1 occupies bit locations 63:32, and so on) and a vector operation that operates on 64-bit data elements (data element position 0 occupies bit locations 63:0, data element position 1 occupies bit locations 127:64, and so on).

Additionally, there is a correlation between the number of one active bit vector writemask elements and the vector size and the data element size according to one embodiment of the invention as shown in FIG. 5. Vector sizes of 128-bits, 256-bits, and 512-bits are shown, although other widths are also possible. Data element sizes of 8-bit bytes (B), 16-bit words (W), 32-bit doublewords (D) or single precision floating point, and 64-bit quadwords (Q) or double precision floating point are considered, although other widths are also possible. As shown, when the vector size is 128-bits, 16-bits may be used for masking when the vector’s data element size is 8-bits, 8-bits may be used for masking when the vector’s data element size is 16-bits, 4-bits may be used for masking when the vector’s data element size is 32-bits, and 2-bits may be used for masking when the vector’s data element size is 64-bits. When the vector size is 256-bits, 32-bits may be used for masking when the packed data element width is 8-bits, 16-bits may be used for masking when the vector’s data element size is 16-bits, 8-bits may be used for masking when the vector’s data element size is 32-bits, and 4-bits may be used for masking when the vector’s data element size is 64-bits. When the vector size is 512-bits, 64-bits may be used for masking when the vector’s data element size is 8-bits, 32-bits may be used for masking when the vector’s data element size is 16-bits, 16-bits may be used for masking when the vector’s data element size is 32-bits, and 8-bits may be used for masking when the vector’s data element size is 64-bits.

Depending upon the combination of the vector size and the data element size, either all 64-bits, or only a subset of the 64-bits, may be used as a write mask. Generally, when a single, per-element masking control bit is used, the number of bits in the vector writemask register used for masking (active bits) is equal to the vector size in bits divided by the vector’s data element size in bits.

A vector instruction that has the flexibility to perform generate couples of subtractions and addition of the vector elements is detailed below. These are fundamental building block for butterfly implementation of transformations like HAAR and HADMARD, which is useful in the field of media codecs and analysis. Below are embodiments of an instruction generically called a packed butterfly, horizontal and cross and or subtract (“PHXADDSUB”) instruction and embodiments of systems, architectures, instruction formats, etc. that may be used to execute such an instruction that is beneficial in several different areas. The execution of a PHXADDSUB instruction causes a horizontal and cross addition or subtraction between packed data elements of a source vector register where the results of the additions and/or subtractions are stored in a destination vector register and the decision of whether to add or subtract is dependent upon a bit in an immediate value of the instruction. More particularly, the source vector register contains one or more data lanes and the horizontal and cross addition or subtraction occurs on a per data lane basis.

The data stored into the destination register is as follows (note that ? is shorthand for a conditional operator which when true will return the first result and when false returns the second result):

For the least significant data element position, the data stored is the result of the least significant data element of the source register being added to or subtracted from the second most least significant data element of the source register wherein the decision of addition or subtraction is based on the least significant bit of the immediate (i.e., $DEST[0]=SRC[0] * (imm8[0] ? -1:1) + SRC[1]$);

For the second most least significant data element position, the data stored is the result of the third most least significant data element of the source register being added to or subtracted from the fourth most least significant data element of the source register wherein the decision of addition or subtraction is based on the third most least significant bit of the immediate ($DEST[1]=SRC[2] * (imm8[2] ? -1:1) + SRC[3]$);

For the third most least significant data element position, the data stored is the result of the second most least significant data element of the source register being added to or subtracted from the least significant data element of the source register wherein the decision of addition or subtraction is based on the second most least significant bit of the immediate ($DEST[2]=SRC[1] * (imm8[1] ? -1:1) + SRC[0]$); and

For the fourth most significant data element position, the data stored is the result of the fourth most least significant data element of the source register being added to or subtracted from the third most least significant data element of the source register wherein the decision of addition or subtraction is based on the fourth most least significant bit of the immediate ($DEST[3]=SRC[3] * (imm8[3] ? -1:1) + SRC[2]$).

FIG. 1 illustrates an exemplary illustration of an operation of an exemplary PHXADDSUB instruction. More particularly, the figure illustrates the operation of this instruction for a source data lane 101 consisting of four packed data

elements into a destination data lane **105**. The calculations performed in this data lane would be done on all of the data lanes of the source register to complete the execution of the instruction.

The source data lane **101** of this example has four packed data elements with the least significant (SRC[0]) being the one on the left. The size of the packed data elements may be many different sizes including 8-bit, 16-bit, 32-bit, 64-bit, etc. The number of data lanes to evaluate is dependent on the size of the data elements and source register size (i.e., 128-, 256-, 512-bit, etc.). For example, with a 128-bit source register and 16-bit data elements, there would be two 64-bit data lanes that would be evaluated as illustrated.

As discussed above, the values to be stored into the destination are as follows:

$$DEST[0]=SRC[0]*(imm8[0]?-1:1)+SRC[1];$$

$$DEST[1]=SRC[2]*(imm8[2]?-1:1)+SRC[3];$$

$$DEST[2]=SRC[1]*(imm8[1]?-1:1)+SRC[0]; \text{ and}$$

$$DEST[3]=SRC[3]*(imm8[3]?-1:1)+SRC[2].$$

In this particular example, only the lower four bits of the immediate **107** are used for the addition/subtraction determination using negation logic **103**. However, if there more data elements, then the number of bits to use in the immediate would change. Additionally, in some embodiments, the upper four bits of the immediate are used.

By way of example, in this figure if the source data element **0** has a value of **3**, the source data element **1** has a value of **2**, and the immediate's least significant bit is **0**, then the resulting value to be stored in destination data element position **0** is 3+2 which equals 5.

Additionally, this figure illustrates addition logic **109** which may be hardware (such as an ALU) or software for performing the addition of the source values (negated or not).

Exemplary Format

An exemplary format of this instruction is "PHXADDSUB{B/W/D/Q} YMM1, YMM2, IMM8" where the operand YMM1 is a destination vector register and YMM2 is a source vector register (such as a 128-, 256-, 512-bit register, etc.), IMM8 is an 8-bit immediate and PHXADDSUB is the instruction's opcode. The size of the data elements may be defined in the "prefix" of the instruction such as through the use of an indication of data granularity bit. In most embodiments, this bit will indicate that each data element is either 32 or 64 bits, however, other variations may be used. However, in some embodiments, the opcode itself will determine the data element size. For example, as shown above, an additional letter {B/W/D/Q} may be used to indicate data element sizes of byte, word, doubleword, and quadword respectively.

Exemplary Methods of Execution

FIG. 2 illustrates an embodiment of the use of a PHXADDSUB instruction in a processor. A PHXADDSUB instruction with a source register operand, destination register operand, and an immediate is fetched at **201**.

The PHXADDSUB instruction is decoded by decoding logic at **203**. Depending on the instruction's format, a variety of data may be interpreted at this stage such as if there is to be a data transformation, which registers to write to and retrieve, what memory address to access, etc.

The source operand values are retrieved/read at **205**. For example, the source register is read. If the source operand is a memory operand, then the data elements associated with

that operand are retrieved. In some embodiments, data elements from memory are stored into a temporary register prior to the execution stage. This stage may also include logically organizing the source register into a plurality of data lanes, wherein each data lane's size is the size of the data elements of the destination register.

The PHXADDSUB instruction (or operations comprising such an instruction such as microoperations) is executed by execution resources such as one or more functional units at **207** to calculate, on a per data lane basis, a horizontal and cross addition or subtraction between packed data elements of a source vector register where the decision of whether to add or subtract is dependent upon a bit in an immediate value of the instruction.

In one embodiment, for the least significant data element position of each data lane, the calculation is the result of the least significant data element of the source register being added to or subtracted from the second most least significant data element of the source register wherein the decision of addition or subtraction is based on the least significant bit of the immediate. For the second most least significant data element position of each data lane, calculation is the result of the third most least significant data element of the source register being added to or subtracted from the fourth most least significant data element of the source register wherein the decision of addition or subtraction is based on the third most least significant bit of the immediate. For the third most least significant data element position of each data lane, the calculation is the result of the second most least significant data element of the source register being added to or subtracted from the least significant data element of the source register wherein the decision of addition or subtraction is based on the second most least significant bit of the immediate. Finally, for the fourth most significant data element position of each data lane, the calculation is the result of the fourth most least significant data element of the source register being added to or subtracted from the third most least significant data element of the source register wherein the decision of addition or subtraction is based on the fourth most least significant bit of the immediate. These calculations may be performed serially or in parallel to each other. For all of the above, when the immediate bit is a **1**, then subtraction is performed and when the immediate bit is a **0** an addition is performed. Of course, in some embodiments, the opposite convention is used.

Each result is stored in a corresponding data element of the destination register at **209**. While **207** and **209** have been illustrated separately, in some embodiments they are performed together as a part of the execution of the instruction.

FIG. 3 illustrates an embodiment of a method for processing a PHXADDSUB instruction. More particularly, this illustration details steps for processing a data lane. For each additional data lane, the same steps may be performed serially or in parallel with other data lanes. In this embodiment it is assumed that some, if not all, of the operations **201-205** have been performed earlier, however, they are not shown in order to not obscure the details presented below. For example, the fetching and decoding are not shown, nor is the operand retrieval shown.

At **301**, the source register is organized into data lanes based data element size and destination register size. As noted above, the number of data elements changes depending on their size and the size of the register that they are in or are going to. However, the typical implementation is to have four data elements in a data lane, with the number of data lanes being variable.

At **303**, a determination, for each data element position of a data lane of the source register, whether the immediate bit associated with that data element position indicates if it should be negated for some of the additions.

A cross addition between adjacent data element positions is performed at **305**. The details of these additions are discussed above.

At **307**, the results are stored into data element positions of the destination register as discussed above.

FIG. 4 illustrates exemplary pseudo-code for a four packed data element destination horizontal add or subtract. Of course, changes to the code would be made to accommodate different sized elements, etc. The condition operator ? means that if true the first result (−1) would return and if false the second result (1) would return.

Exemplary Instruction Formats

Embodiments of the instruction(s) described herein may be embodied in different formats. For example, the instruction(s) described herein may be embodied as a VEX, generic vector friendly, or other format. Details of VEX and a generic vector friendly format are discussed below. Additionally, exemplary systems, architectures, and pipelines are detailed below. Embodiments of the instruction(s) may be executed on such systems, architectures, and pipelines, but are not limited to those detailed.

VEX Instruction Format

VEX encoding allows instructions to have more than two operands, and allows SIMD vector registers to be longer than 128 bits. The use of a VEX prefix provides for three-operand (or more) syntax. For example, previous two-operand instructions performed operations such as $A=A+B$, which overwrites a source operand. The use of a VEX prefix enables operands to perform nondestructive operations such as $A=B+C$.

FIG. 6A illustrates an exemplary AVX instruction format including a VEX prefix **602**, real opcode field **630**, Mod R/M byte **640**, SIB byte **650**, displacement field **662**, and IMM8 **672**. FIG. 6B illustrates which fields from FIG. 6A make up a full opcode field **674** and a base operation field **642**. FIG. 6C illustrates which fields from FIG. 6A make up a register index field **644**.

VEX Prefix (Bytes **0-2**) **602** is encoded in a three-byte form. The first byte is the Format Field **640** (VEX Byte **0**, bits [7:0]), which contains an explicit C4 byte value (the unique value used for distinguishing the C4 instruction format). The second-third bytes (VEX Bytes **1-2**) include a number of bit fields providing specific capability. Specifically, REX field **605** (VEX Byte **1**, bits [7:5]) consists of a VEX.R bit field (VEX Byte **1**, bit [7]-R), VEX.X bit field (VEX byte **1**, bit [6]-X), and VEX.B bit field (VEX byte **1**, bit [5]-B). Other fields of the instructions encode the lower three bits of the register indexes as is known in the art (rrr, xxx, and bbb), so that Rrrr, Xxxx, and Bbbb may be formed by adding VEX.R, VEX.X, and VEX.B. Opcode map field **615** (VEX byte **1**, bits [4:0]-mmmmm) includes content to encode an implied leading opcode byte. W Field **664** (VEX byte **2**, bit [7]-W)—is represented by the notation VEX.W, and provides different functions depending on the instruction. The role of VEX.vvvv **620** (VEX Byte **2**, bits [6:3]-vvvv) may include the following: 1) VEX.vvvv encodes the first source register operand, specified in inverted (1's complement) form and is valid for instructions with 2 or more source operands; 2) VEX.vvvv encodes the destination register operand, specified in is complement form for certain vector shifts; or 3) VEX.vvvv does not encode any operand, the field is reserved and should contain 1111b. If VEX.L **668** Size field (VEX byte **2**, bit [2]-L)=0, it indicates 128 bit

vector; if VEX.L=1, it indicates 256 bit vector. Prefix encoding field **625** (VEX byte **2**, bits [1:0]-pp) provides additional bits for the base operation field.

Real Opcode Field **630** (Byte **3**) is also known as the opcode byte. Part of the opcode is specified in this field.

MOD R/M Field **640** (Byte **4**) includes MOD field **642** (bits [7-6]), Reg field **644** (bits [5-3]), and R/M field **646** (bits [2-0]). The role of Reg field **644** may include the following: encoding either the destination register operand or a source register operand (the rrr of Rrrr), or be treated as an opcode extension and not used to encode any instruction operand. The role of R/M field **646** may include the following: encoding the instruction operand that references a memory address, or encoding either the destination register operand or a source register operand.

Scale, Index, Base (SIB)—The content of Scale field **650** (Byte **5**) includes SS**652** (bits [7-6]), which is used for memory address generation. The contents of SIB.xxx **654** (bits [5-3]) and SIB.bbb **656** (bits [2-0]) have been previously referred to with regard to the register indexes Xxxx and Bbbb.

The Displacement Field **662** and the immediate field (IMM8) **672** contain address data.

Generic Vector Friendly Instruction Format

A vector friendly instruction format is an instruction format that is suited for vector instructions (e.g., there are certain fields specific to vector operations). While embodiments are described in which both vector and scalar operations are supported through the vector friendly instruction format, alternative embodiments use only vector operations the vector friendly instruction format.

FIGS. 7A-7B are block diagrams illustrating a generic vector friendly instruction format and instruction templates thereof according to embodiments of the invention. FIG. 7A is a block diagram illustrating a generic vector friendly instruction format and class A instruction templates thereof according to embodiments of the invention; while FIG. 7B is a block diagram illustrating the generic vector friendly instruction format and class B instruction templates thereof according to embodiments of the invention. Specifically, a generic vector friendly instruction format **700** for which are defined class A and class B instruction templates, both of which include no memory access **705** instruction templates and memory access **720** instruction templates. The term generic in the context of the vector friendly instruction format refers to the instruction format not being tied to any specific instruction set.

While embodiments of the invention will be described in which the vector friendly instruction format supports the following: a 64 byte vector operand length (or size) with 32 bit (4 byte) or 64 bit (8 byte) data element widths (or sizes) (and thus, a 64 byte vector consists of either 16 doubleword-size elements or alternatively, 8 quadword-size elements); a 64 byte vector operand length (or size) with 16 bit (2 byte) or 8 bit (1 byte) data element widths (or sizes); a 32 byte vector operand length (or size) with 32 bit (4 byte), 64 bit (8 byte), 16 bit (2 byte), or 8 bit (1 byte) data element widths (or sizes); and a 16 byte vector operand length (or size) with 32 bit (4 byte), 64 bit (8 byte), 16 bit (2 byte), or 8 bit (1 byte) data element widths (or sizes); alternative embodiments may support more, less and/or different vector operand and sizes (e.g., 256 byte vector operands) with more, less, or different data element widths (e.g., 128 bit (16 byte) data element widths).

The class A instruction templates in FIG. 7A include: 1) within the no memory access **705** instruction templates there is shown a no memory access, full round control type

11

operation **710** instruction template and a no memory access, data transform type operation **715** instruction template; and 2) within the memory access **720** instruction templates there is shown a memory access, temporal **725** instruction template and a memory access, non-temporal **730** instruction template. The class B instruction templates in FIG. 7B include: 1) within the no memory access **705** instruction templates there is shown a no memory access, write mask control, partial round control type operation **712** instruction template and a no memory access, write mask control, vsize type operation **717** instruction template; and 2) within the memory access **720** instruction templates there is shown a memory access, write mask control **727** instruction template.

The generic vector friendly instruction format **700** includes the following fields listed below in the order illustrated in FIGS. 7A-7B.

Format field **740**—a specific value (an instruction format identifier value) in this field uniquely identifies the vector friendly instruction format, and thus occurrences of instructions in the vector friendly instruction format in instruction streams. As such, this field is optional in the sense that it is not needed for an instruction set that has only the generic vector friendly instruction format.

Base operation field **742**—its content distinguishes different base operations.

Register index field **744**—its content, directly or through address generation, specifies the locations of the source and destination operands, be they in registers or in memory. These include a sufficient number of bits to select N registers from a P×Q (e.g. 32×512, 16×128, 32×1024, 64×1024) register file. While in one embodiment N may be up to three sources and one destination register, alternative embodiments may support more or less sources and destination registers (e.g., may support up to two sources where one of these sources also acts as the destination, may support up to three sources where one of these sources also acts as the destination, may support up to two sources and one destination).

Modifier field **746**—its content distinguishes occurrences of instructions in the generic vector instruction format that specify memory access from those that do not; that is, between no memory access **705** instruction templates and memory access **720** instruction templates. Memory access operations read and/or write to the memory hierarchy (in some cases specifying the source and/or destination addresses using values in registers), while non-memory access operations do not (e.g., the source and destinations are registers). While in one embodiment this field also selects between three different ways to perform memory address calculations, alternative embodiments may support more, less, or different ways to perform memory address calculations.

Augmentation operation field **750**—its content distinguishes which one of a variety of different operations to be performed in addition to the base operation. This field is context specific. In one embodiment of the invention, this field is divided into a class field **768**, an alpha field **752**, and a beta field **754**. The augmentation operation field **750** allows common groups of operations to be performed in a single instruction rather than 2, 3, or 4 instructions.

Scale field **760**—its content allows for the scaling of the index field's content for memory address generation (e.g., for address generation that uses $2^{scale} \times \text{index} + \text{base}$).

Displacement Field **762A**—its content is used as part of memory address generation (e.g., for address generation that uses $2^{scale} \times \text{index} + \text{base} + \text{displacement}$).

12

Displacement Factor Field **762B** (note that the juxtaposition of displacement field **762A** directly over displacement factor field **762B** indicates one or the other is used)—its content is used as part of address generation; it specifies a displacement factor that is to be scaled by the size of a memory access (N)—where N is the number of bytes in the memory access (e.g., for address generation that uses $2^{scale} \times \text{index} + \text{base} + \text{scaled displacement}$). Redundant low-order bits are ignored and hence, the displacement factor field's content is multiplied by the memory operands total size (N) in order to generate the final displacement to be used in calculating an effective address. The value of N is determined by the processor hardware at runtime based on the full opcode field **774** (described later herein) and the data manipulation field **754C**. The displacement field **762A** and the displacement factor field **762B** are optional in the sense that they are not used for the no memory access **705** instruction templates and/or different embodiments may implement only one or none of the two.

Data element width field **764**—its content distinguishes which one of a number of data element widths is to be used (in some embodiments for all instructions; in other embodiments for only some of the instructions). This field is optional in the sense that it is not needed if only one data element width is supported and/or data element widths are supported using some aspect of the opcodes.

Write mask field **770**—its content controls, on a per data element position basis, whether that data element position in the destination vector operand reflects the result of the base operation and augmentation operation. Class A instruction templates support merging-writemasking, while class B instruction templates support both merging- and zeroing-writemasking. When merging, vector masks allow any set of elements in the destination to be protected from updates during the execution of any operation (specified by the base operation and the augmentation operation); in other one embodiment, preserving the old value of each element of the destination where the corresponding mask bit has a 0. In contrast, when zeroing vector masks allow any set of elements in the destination to be zeroed during the execution of any operation (specified by the base operation and the augmentation operation); in one embodiment, an element of the destination is set to 0 when the corresponding mask bit has a 0 value. A subset of this functionality is the ability to control the vector length of the operation being performed (that is, the span of elements being modified, from the first to the last one); however, it is not necessary that the elements that are modified be consecutive. Thus, the write mask field **770** allows for partial vector operations, including loads, stores, arithmetic, logical, etc. While embodiments of the invention are described in which the write mask field's **770** content selects one of a number of write mask registers that contains the write mask to be used (and thus the write mask field's **770** content indirectly identifies that masking to be performed), alternative embodiments instead or additionally allow the mask write field's **770** content to directly specify the masking to be performed.

Immediate field **772**—its content allows for the specification of an immediate. This field is optional in the sense that it is not present in an implementation of the generic vector friendly format that does not support immediate and it is not present in instructions that do not use an immediate.

Class field **768**—its content distinguishes between different classes of instructions. With reference to FIGS. 7A-B, the contents of this field select between class A and class B instructions. In FIGS. 7A-B, rounded corner squares are

13

used to indicate a specific value is present in a field (e.g., class A **768A** and class B **768B** for the class field **768** respectively in FIGS. 7A-B).

Instruction Templates of Class A

In the case of the non-memory access **705** instruction templates of class A, the alpha field **752** is interpreted as an RS field **752A**, whose content distinguishes which one of the different augmentation operation types are to be performed (e.g., round **752A.1** and data transform **752A.2** are respectively specified for the no memory access, round type operation **710** and the no memory access, data transform type operation **715** instruction templates), while the beta field **754** distinguishes which of the operations of the specified type is to be performed. In the no memory access **705** instruction templates, the scale field **760**, the displacement field **762A**, and the displacement scale field **762B** are not present.

No-Memory Access Instruction Templates—Full Round Control Type Operation

In the no memory access full round control type operation **710** instruction template, the beta field **754** is interpreted as a round control field **754A**, whose content(s) provide static rounding. While in the described embodiments of the invention the round control field **754A** includes a suppress all floating point exceptions (SAE) field **756** and a round operation control field **758**, alternative embodiments may support may encode both these concepts into the same field or only have one or the other of these concepts/fields (e.g., may have only the round operation control field **758**).

SAE field **756**—its content distinguishes whether or not to disable the exception event reporting; when the SAE field's **756** content indicates suppression is enabled, a given instruction does not report any kind of floating-point exception flag and does not raise any floating point exception handler.

Round operation control field **758**—its content distinguishes which one of a group of rounding operations to perform (e.g., Round-up, Round-down, Round-towards-zero and Round-to-nearest). Thus, the round operation control field **758** allows for the changing of the rounding mode on a per instruction basis. In one embodiment of the invention where a processor includes a control register for specifying rounding modes, the round operation control field's **750** content overrides that register value.

No Memory Access Instruction Templates—Data Transform Type Operation

In the no memory access data transform type operation **715** instruction template, the beta field **754** is interpreted as a data transform field **754B**, whose content distinguishes which one of a number of data transforms is to be performed (e.g., no data transform, swizzle, broadcast).

In the case of a memory access **720** instruction template of class A, the alpha field **752** is interpreted as an eviction hint field **752B**, whose content distinguishes which one of the eviction hints is to be used (in FIG. 7A, temporal **752B.1** and non-temporal **752B.2** are respectively specified for the memory access, temporal **725** instruction template and the memory access, non-temporal **730** instruction template), while the beta field **754** is interpreted as a data manipulation field **754C**, whose content distinguishes which one of a number of data manipulation operations (also known as primitives) is to be performed (e.g., no manipulation; broadcast; up conversion of a source; and down conversion of a destination). The memory access **720** instruction templates include the scale field **760**, and optionally the displacement field **762A** or the displacement scale field **762B**.

14

Vector memory instructions perform vector loads from and vector stores to memory, with conversion support. As with regular vector instructions, vector memory instructions transfer data from/to memory in a data element-wise fashion, with the elements that are actually transferred is dictated by the contents of the vector mask that is selected as the write mask.

Memory Access Instruction Templates—Temporal

Temporal data is data likely to be reused soon enough to benefit from caching. This is, however, a hint, and different processors may implement it in different ways, including ignoring the hint entirely.

Memory Access Instruction Templates—Non-Temporal

Non-temporal data is data unlikely to be reused soon enough to benefit from caching in the 1st-level cache and should be given priority for eviction. This is, however, a hint, and different processors may implement it in different ways, including ignoring the hint entirely.

Instruction Templates of Class B

In the case of the instruction templates of class B, the alpha field **752** is interpreted as a write mask control (Z) field **752C**, whose content distinguishes whether the write masking controlled by the write mask field **770** should be a merging or a zeroing.

In the case of the non-memory access **705** instruction templates of class B, part of the beta field **754** is interpreted as an RL field **757A**, whose content distinguishes which one of the different augmentation operation types are to be performed (e.g., round **757A.1** and vector length (VSIZE) **757A.2** are respectively specified for the no memory access, write mask control, partial round control type operation **712** instruction template and the no memory access, write mask control, VSIZE type operation **717** instruction template), while the rest of the beta field **754** distinguishes which of the operations of the specified type is to be performed. In the no memory access **705** instruction templates, the scale field **760**, the displacement field **762A**, and the displacement scale field **762B** are not present.

In the no memory access, write mask control, partial round control type operation **710** instruction template, the rest of the beta field **754** is interpreted as a round operation field **759A** and exception event reporting is disabled (a given instruction does not report any kind of floating-point exception flag and does not raise any floating point exception handler).

Round operation control field **759A**—just as round operation control field **758**, its content distinguishes which one of a group of rounding operations to perform (e.g., Round-up, Round-down, Round-towards-zero and Round-to-nearest). Thus, the round operation control field **759A** allows for the changing of the rounding mode on a per instruction basis. In one embodiment of the invention where a processor includes a control register for specifying rounding modes, the round operation control field's **750** content overrides that register value.

In the no memory access, write mask control, VSIZE type operation **717** instruction template, the rest of the beta field **754** is interpreted as a vector length field **759B**, whose content distinguishes which one of a number of data vector lengths is to be performed on (e.g., 128, 256, or 512 byte).

In the case of a memory access **720** instruction template of class B, part of the beta field **754** is interpreted as a broadcast field **757B**, whose content distinguishes whether or not the broadcast type data manipulation operation is to be performed, while the rest of the beta field **754** is interpreted the vector length field **759B**. The memory access **720**

15

instruction templates include the scale field **760**, and optionally the displacement field **762A** or the displacement scale field **762B**.

With regard to the generic vector friendly instruction format **700**, a full opcode field **774** is shown including the format field **740**, the base operation field **742**, and the data element width field **764**. While one embodiment is shown where the full opcode field **774** includes all of these fields, the full opcode field **774** includes less than all of these fields in embodiments that do not support all of them. The full opcode field **774** provides the operation code (opcode).

The augmentation operation field **750**, the data element width field **764**, and the write mask field **770** allow these features to be specified on a per instruction basis in the generic vector friendly instruction format.

The combination of write mask field and data element width field create typed instructions in that they allow the mask to be applied based on different data element widths.

The various instruction templates found within class A and class B are beneficial in different situations. In some embodiments of the invention, different processors or different cores within a processor may support only class A, only class B, or both classes. For instance, a high performance general purpose out-of-order core intended for general-purpose computing may support only class B, a core intended primarily for graphics and/or scientific (throughput) computing may support only class A, and a core intended for both may support both (of course, a core that has some mix of templates and instructions from both classes but not all templates and instructions from both classes is within the purview of the invention). Also, a single processor may include multiple cores, all of which support the same class or in which different cores support different class. For instance, in a processor with separate graphics and general purpose cores, one of the graphics cores intended primarily for graphics and/or scientific computing may support only class A, while one or more of the general purpose cores may be high performance general purpose cores with out of order execution and register renaming intended for general-purpose computing that support only class B. Another processor that does not have a separate graphics core, may include one more general purpose in-order or out-of-order cores that support both class A and class B. Of course, features from one class may also be implemented in the other class in different embodiments of the invention. Programs written in a high level language would be put (e.g., just in time compiled or statically compiled) into a variety of different executable forms, including: 1) a form having only instructions of the class(es) supported by the target processor for execution; or 2) a form having alternative routines written using different combinations of the instructions of all classes and having control flow code that selects the routines to execute based on the instructions supported by the processor which is currently executing the code.

Exemplary Specific Vector Friendly Instruction Format

FIG. 8 is a block diagram illustrating an exemplary specific vector friendly instruction format according to embodiments of the invention. FIG. 8 shows a specific vector friendly instruction format **800** that is specific in the sense that it specifies the location, size, interpretation, and order of the fields, as well as values for some of those fields. The specific vector friendly instruction format **800** may be used to extend the x86 instruction set, and thus some of the fields are similar or the same as those used in the existing x86 instruction set and extension thereof (e.g., AVX). This format remains consistent with the prefix encoding field, real

16

opcode byte field, MOD R/M field, SIB field, displacement field, and immediate fields of the existing x86 instruction set with extensions. The fields from FIG. 7 into which the fields from FIG. 8 map are illustrated.

It should be understood that, although embodiments of the invention are described with reference to the specific vector friendly instruction format **800** in the context of the generic vector friendly instruction format **700** for illustrative purposes, the invention is not limited to the specific vector friendly instruction format **800** except where claimed. For example, the generic vector friendly instruction format **700** contemplates a variety of possible sizes for the various fields, while the specific vector friendly instruction format **800** is shown as having fields of specific sizes. By way of specific example, while the data element width field **764** is illustrated as a one bit field in the specific vector friendly instruction format **800**, the invention is not so limited (that is, the generic vector friendly instruction format **700** contemplates other sizes of the data element width field **764**).

The generic vector friendly instruction format **700** includes the following fields listed below in the order illustrated in FIG. 8A.

EVEX Prefix (Bytes 0-3) **802**—is encoded in a four-byte form.

Format Field **740** (EVEX Byte 0, bits [7:0])—the first byte (EVEX Byte 0) is the format field **740** and it contains 0x62 (the unique value used for distinguishing the vector friendly instruction format in one embodiment of the invention).

The second-fourth bytes (EVEX Bytes 1-3) include a number of bit fields providing specific capability.

REX field **805** (EVEX Byte 1, bits [7:5])—consists of a EVEX.R bit field (EVEX Byte 1, bit [7]-R), EVEX.X bit field (EVEX byte 1, bit [6]-X), and **757BEX** byte 1, bit[5]-B). The EVEX.R, EVEX.X, and EVEX.B bit fields provide the same functionality as the corresponding VEX bit fields, and are encoded using is complement form, i.e. ZMM0 is encoded as 1111B, ZMM15 is encoded as 0000B. Other fields of the instructions encode the lower three bits of the register indexes as is known in the art (rrr, xxx, and bbb), so that Rrrr, Xxxx, and Bbbb may be formed by adding EVEX.R, EVEX.X, and EVEX.B.

REX' field **710**—this is the first part of the REX' field **710** and is the EVEX.R' bit field (EVEX Byte 1, bit [4]-R') that is used to encode either the upper **16** or lower **16** of the extended **32** register set. In one embodiment of the invention, this bit, along with others as indicated below, is stored in bit inverted format to distinguish (in the well-known x86 32-bit mode) from the BOUND instruction, whose real opcode byte is **62**, but does not accept in the MOD R/M field (described below) the value of **11** in the MOD field; alternative embodiments of the invention do not store this and the other indicated bits below in the inverted format. A value of **1** is used to encode the lower 16 registers. In other words, R'Rrrr is formed by combining EVEX.R', EVEX.R, and the other RRR from other fields.

Opcode map field **815** (EVEX byte 1, bits [3:0]-mmmm)—its content encodes an implied leading opcode byte (0F, 0F 38, or 0F 3).

Data element width field **764** (EVEX byte 2, bit [7]-W)—is represented by the notation EVEX.W. EVEX.W is used to define the granularity (size) of the datatype (either 32-bit data elements or 64-bit data elements).

EVEX.vvvv **820** (EVEX Byte 2, bits [6:3]-vvvv)—the role of EVEX.vvvv may include the following: 1) EVEX.vvvv encodes the first source register operand, specified in inverted (1's complement) form and is valid for instructions

with 2 or more source operands; 2) EVEX.vvvv encodes the destination register operand, specified in its complement form for certain vector shifts; or 3) EVEX.vvvv does not encode any operand, the field is reserved and should contain 1111b. Thus, EVEX.vvvv field **820** encodes the 4 low-order bits of the first source register specifier stored in inverted (1's complement) form. Depending on the instruction, an extra different EVEX bit field is used to extend the specifier size to 32 registers.

EVEX.U **768** Class field (EVEX byte 2, bit [2]-U)—If EVEX.0=0, it indicates class A or EVEX.U0; if EVEX.0=1, it indicates class B or EVEX.U1.

Prefix encoding field **825** (EVEX byte 2, bits [1:0]-pp)—provides additional bits for the base operation field. In addition to providing support for the legacy SSE instructions in the EVEX prefix format, this also has the benefit of compacting the SIMD prefix (rather than requiring a byte to express the SIMD prefix, the EVEX prefix requires only 2 bits). In one embodiment, to support legacy SSE instructions that use a SIMD prefix (66H, F2H, F3H) in both the legacy format and in the EVEX prefix format, these legacy SIMD prefixes are encoded into the SIMD prefix encoding field; and at runtime are expanded into the legacy SIMD prefix prior to being provided to the decoder's PLA (so the PLA can execute both the legacy and EVEX format of these legacy instructions without modification). Although newer instructions could use the EVEX prefix encoding field's content directly as an opcode extension, certain embodiments expand in a similar fashion for consistency but allow for different meanings to be specified by these legacy SIMD prefixes. An alternative embodiment may redesign the PLA to support the 2 bit SIMD prefix encodings, and thus not require the expansion.

Alpha field **752** (EVEX byte 3, bit [7]-EH; also known as EVEX.EH, EVEX.rs, EVEX.RL, EVEX.write mask control, and EVEX.N; also illustrated with α)—as previously described, this field is context specific.

Beta field **754** (EVEX byte 3, bits [6:4]-SSS, also known as EVEX.s₂₋₀, EVEX.r₂₋₀, EVEX.rr1, EVEX.LL0, EVEX.LLB; also illustrated with $\beta\beta\beta$)—as previously described, this field is context specific.

REX' field **710**—this is the remainder of the REX' field and is the EVEX.V' bit field (EVEX Byte 3, bit [3]-V') that may be used to encode either the upper 16 or lower 16 of the extended 32 register set. This bit is stored in bit inverted format. A value of 1 is used to encode the lower 16 registers. In other words, V'VVVV is formed by combining EVEX.V', EVEX.vvvv.

Write mask field **770** (EVEX byte 3, bits [2:0]-kkk)—its content specifies the index of a register in the write mask registers as previously described. In one embodiment of the invention, the specific value EVEX.kkk=000 has a special behavior implying no write mask is used for the particular instruction (this may be implemented in a variety of ways including the use of a write mask hardwired to all ones or hardware that bypasses the masking hardware).

Real Opcode Field **830** (Byte 4) is also known as the opcode byte. Part of the opcode is specified in this field.

MOD R/M Field **840** (Byte 5) includes MOD field **842**, Reg field **844**, and R/M field **846**. As previously described, the MOD field's **842** content distinguishes between memory access and non-memory access operations. The role of Reg field **844** can be summarized to two situations: encoding either the destination register operand or a source register operand, or be treated as an opcode extension and not used to encode any instruction operand. The role of R/M field **846** may include the following: encoding the instruction operand

that references a memory address, or encoding either the destination register operand or a source register operand.

Scale, Index, Base (SIB) Byte (Byte 6)—As previously described, the scale field's **750** content is used for memory address generation. SIB.xxx **854** and SIB.bbb **856**—the contents of these fields have been previously referred to with regard to the register indexes Xxxx and Bbbb.

Displacement field **762A** (Bytes 7-10)—when MOD field **842** contains 10, bytes 7-10 are the displacement field **762A**, and it works the same as the legacy 32-bit displacement (disp32) and works at byte granularity.

Displacement factor field **762B** (Byte 7)—when MOD field **842** contains 01, byte 7 is the displacement factor field **762B**. The location of this field is that same as that of the legacy x86 instruction set 8-bit displacement (disp8), which works at byte granularity. Since disp8 is sign extended, it can only address between -128 and 127 bytes offsets; in terms of 64 byte cache lines, disp8 uses 8 bits that can be set to only four really useful values -128, -64, 0, and 64; since a greater range is often needed, disp32 is used; however, disp32 requires 4 bytes. In contrast to disp8 and disp32, the displacement factor field **762B** is a reinterpretation of disp8; when using displacement factor field **762B**, the actual displacement is determined by the content of the displacement factor field multiplied by the size of the memory operand access (N). This type of displacement is referred to as disp8*N. This reduces the average instruction length (a single byte of used for the displacement but with a much greater range). Such compressed displacement is based on the assumption that the effective displacement is multiple of the granularity of the memory access, and hence, the redundant low-order bits of the address offset do not need to be encoded. In other words, the displacement factor field **762B** substitutes the legacy x86 instruction set 8-bit displacement. Thus, the displacement factor field **762B** is encoded the same way as an x86 instruction set 8-bit displacement (so no changes in the ModRM/SIB encoding rules) with the only exception that disp8 is overloaded to disp8*N. In other words, there are no changes in the encoding rules or encoding lengths but only in the interpretation of the displacement value by hardware (which needs to scale the displacement by the size of the memory operand to obtain a byte-wise address offset).

Immediate field **772** operates as previously described.

Full Opcode Field

FIG. 8B is a block diagram illustrating the fields of the specific vector friendly instruction format **800** that make up the full opcode field **774** according to one embodiment of the invention. Specifically, the full opcode field **774** includes the format field **740**, the base operation field **742**, and the data element width (W) field **764**. The base operation field **742** includes the prefix encoding field **825**, the opcode map field **815**, and the real opcode field **830**.

Register Index Field

FIG. 8C is a block diagram illustrating the fields of the specific vector friendly instruction format **800** that make up the register index field **744** according to one embodiment of the invention. Specifically, the register index field **744** includes the REX field **805**, the REX' field **810**, the MODR/M.reg field **844**, the MODR/M.r/m field **846**, the VVVV field **820**, xxx field **854**, and the bbb field **856**.

Augmentation Operation Field

FIG. 8D is a block diagram illustrating the fields of the specific vector friendly instruction format **800** that make up the augmentation operation field **750** according to one embodiment of the invention. When the class (U) field **768** contains 0, it signifies EVEX.U0 (class A **768A**); when it

contains **1**, it signifies EVEX.U1 (class B **768B**). When U=0 and the MOD field **842** contains **11** (signifying a no memory access operation), the alpha field **752** (EVEX byte 3, bit [7]-EH) is interpreted as the rs field **752A**. When the rs field **752A** contains a **1** (round **752A.1**), the beta field **754** (EVEX byte 3, bits [6:4]-SSS) is interpreted as the round control field **754A**. The round control field **754A** includes a one bit SAE field **756** and a two bit round operation field **758**. When the rs field **752A** contains a **0** (data transform **752A.2**), the beta field **754** (EVEX byte 3, bits [6:4]-SSS) is interpreted as a three bit data transform field **754B**. When U=0 and the MOD field **842** contains **00**, **01**, or **10** (signifying a memory access operation), the alpha field **752** (EVEX byte 3, bit [7]-EH) is interpreted as the eviction hint (EH) field **752B** and the beta field **754** (EVEX byte 3, bits [6:4]-SSS) is interpreted as a three bit data manipulation field **754C**.

When U=1, the alpha field **752** (EVEX byte 3, bit [7]-EH) is interpreted as the write mask control (Z) field **752C**. When U=1 and the MOD field **842** contains **11** (signifying a no memory access operation), part of the beta field **754** (EVEX byte 3, bit [4]-S₀) is interpreted as the RL field **757A**; when it contains a **1** (round **757A.1**) the rest of the beta field **754** (EVEX byte 3, bit [6:5]-S₂₋₁) is interpreted as the round operation field **759A**, while when the RL field **757A** contains a **0** (VSIZE **757A.2**) the rest of the beta field **754** (EVEX byte 3, bit [6:5]-S₂₋₁) is interpreted as the vector length field **759B** (EVEX byte 3, bit [6:5]-L₁₋₀). When U=1 and the MOD field **842** contains **00**, **01**, or **10** (signifying a memory access operation), the beta field **754** (EVEX byte 3, bits [6:4]-SSS) is interpreted as the vector length field **759B** (EVEX byte 3, bit [6:5]-L₁₋₀) and the broadcast field **757B** (EVEX byte 3, bit [4]-B).

Exemplary Register Architecture

FIG. 9 is a block diagram of a register architecture **900** according to one embodiment of the invention. In the embodiment illustrated, there are 32 vector registers **910** that are 512 bits wide; these registers are referenced as zmm0 through zmm31. The lower order 256 bits of the lower 16 zmm registers are overlaid on registers ymm0-16. The lower order 128 bits of the lower 16 zmm registers (the lower order 128 bits of the ymm registers) are overlaid on registers xmm0-15. The specific vector friendly instruction format **800** operates on these overlaid register file as illustrated in the below tables.

Adjustable Vector Length	Class	Operations	Registers
Instruction Templates that do not include the vector length field 759B	A (FIG. 7A; U = 0)	710, 715, 725, 730	zmm registers (the vector length is 64 byte)
	B (FIG. 7B; U = 1)	712	zmm registers (the vector length is 64 byte)
Instruction Templates that do include the vector length field 759B	B (FIG. 7B; U = 1)	717, 727	zmm, ymm, or xmm registers (the vector length is 64 byte, 32 byte, or 16 byte) depending on the vector length field 759B

In other words, the vector length field **759B** selects between a maximum length and one or more other shorter lengths, where each such shorter length is half the length of the preceding length; and instructions templates without the vector length field **759B** operate on the maximum vector length. Further, in one embodiment, the class B instruction

templates of the specific vector friendly instruction format **800** operate on packed or scalar single/double-precision floating point data and packed or scalar integer data. Scalar operations are operations performed on the lowest order data element position in an zmm/ymm/xmm register; the higher order data element positions are either left the same as they were prior to the instruction or zeroed depending on the embodiment.

Write mask registers **915**—in the embodiment illustrated, there are 8 write mask registers (k0 through k7), each 64 bits in size. In an alternate embodiment, the write mask registers **915** are 16 bits in size. As previously described, in one embodiment of the invention, the vector mask register k0 cannot be used as a write mask; when the encoding that would normally indicate k0 is used for a write mask, it selects a hardwired write mask of 0xFFFF, effectively disabling write masking for that instruction.

General-purpose registers **925**—in the embodiment illustrated, there are sixteen 64-bit general-purpose registers that are used along with the existing x86 addressing modes to address memory operands. These registers are referenced by the names RAX, RBX, RCX, RDX, RBP, RSI, RDI, RSP, and R8 through R15.

Scalar floating point stack register file (x87 stack) **945**, on which is aliased the MMX packed integer flat register file **950**—in the embodiment illustrated, the x87 stack is an eight-element stack used to perform scalar floating-point operations on 32/64/80-bit floating point data using the x87 instruction set extension; while the MMX registers are used to perform operations on 64-bit packed integer data, as well as to hold operands for some operations performed between the MMX and XMM registers.

Alternative embodiments of the invention may use wider or narrower registers. Additionally, alternative embodiments of the invention may use more, less, or different register files and registers.

Exemplary Core Architectures, Processors, and Computer Architectures

Processor cores may be implemented in different ways, for different purposes, and in different processors. For instance, implementations of such cores may include: 1) a general purpose in-order core intended for general-purpose computing; 2) a high performance general purpose out-of-order core intended for general-purpose computing; 3) a special purpose core intended primarily for graphics and/or scientific (throughput) computing. Implementations of different processors may include: 1) a CPU including one or more general purpose in-order cores intended for general-purpose computing and/or one or more general purpose out-of-order cores intended for general-purpose computing; and 2) a coprocessor including one or more special purpose cores intended primarily for graphics and/or scientific (throughput). Such different processors lead to different computer system architectures, which may include: 1) the coprocessor on a separate chip from the CPU; 2) the coprocessor on a separate die in the same package as a CPU; 3) the coprocessor on the same die as a CPU (in which case, such a coprocessor is sometimes referred to as special purpose logic, such as integrated graphics and/or scientific (throughput) logic, or as special purpose cores); and 4) a system on a chip that may include on the same die the described CPU (sometimes referred to as the application core(s) or application processor(s)), the above described coprocessor, and additional functionality. Exemplary core architectures are described next, followed by descriptions of exemplary processors and computer architectures.

Exemplary Core Architectures

In-order and out-of-order core block diagram

FIG. 10A is a block diagram illustrating both an exemplary in-order pipeline and an exemplary register renaming, out-of-order issue/execution pipeline according to embodiments of the invention. FIG. 10B is a block diagram illustrating both an exemplary embodiment of an in-order architecture core and an exemplary register renaming, out-of-order issue/execution architecture core to be included in a processor according to embodiments of the invention. The solid lined boxes in FIGS. 10A-B illustrate the in-order pipeline and in-order core, while the optional addition of the dashed lined boxes illustrates the register renaming, out-of-order issue/execution pipeline and core. Given that the in-order aspect is a subset of the out-of-order aspect, the out-of-order aspect will be described.

In FIG. 10A, a processor pipeline 1000 includes a fetch stage 1002, a length decode stage 1004, a decode stage 1006, an allocation stage 1008, a renaming stage 1010, a scheduling (also known as a dispatch or issue) stage 1012, a register read/memory read stage 1014, an execute stage 1016, a write back/memory write stage 1018, an exception handling stage 1022, and a commit stage 1024.

FIG. 10B shows processor core 1090 including a front end unit 1030 coupled to an execution engine unit 1050, and both are coupled to a memory unit 1070. The core 1090 may be a reduced instruction set computing (RISC) core, a complex instruction set computing (CISC) core, a very long instruction word (VLIW) core, or a hybrid or alternative core type. As yet another option, the core 1090 may be a special-purpose core, such as, for example, a network or communication core, compression engine, coprocessor core, general purpose computing graphics processing unit (GPGPU) core, graphics core, or the like.

The front end unit 1030 includes a branch prediction unit 1032 coupled to an instruction cache unit 1034, which is coupled to an instruction translation lookaside buffer (TLB) 1036, which is coupled to an instruction fetch unit 1038, which is coupled to a decode unit 1040. The decode unit 1040 (or decoder) may decode instructions, and generate as an output one or more micro-operations, micro-code entry points, microinstructions, other instructions, or other control signals, which are decoded from, or which otherwise reflect, or are derived from, the original instructions. The decode unit 1040 may be implemented using various different mechanisms. Examples of suitable mechanisms include, but are not limited to, look-up tables, hardware implementations, programmable logic arrays (PLAs), microcode read only memories (ROMs), etc. In one embodiment, the core 1090 includes a microcode ROM or other medium that stores microcode for certain macroinstructions (e.g., in decode unit 1040 or otherwise within the front end unit 1030). The decode unit 1040 is coupled to a rename/allocator unit 1052 in the execution engine unit 1050.

The execution engine unit 1050 includes the rename/allocator unit 1052 coupled to a retirement unit 1054 and a set of one or more scheduler unit(s) 1056. The scheduler unit(s) 1056 represents any number of different schedulers, including reservations stations, central instruction window, etc. The scheduler unit(s) 1056 is coupled to the physical register file(s) unit(s) 1058. Each of the physical register file(s) units 1058 represents one or more physical register files, different ones of which store one or more different data types, such as scalar integer, scalar floating point, packed integer, packed floating point, vector integer, vector floating point, status (e.g., an instruction pointer that is the address of the next instruction to be executed), etc. In one embodi-

ment, the physical register file(s) unit 1058 comprises a vector registers unit, a write mask registers unit, and a scalar registers unit. These register units may provide architectural vector registers, vector mask registers, and general purpose registers. The physical register file(s) unit(s) 1058 is overlapped by the retirement unit 1054 to illustrate various ways in which register renaming and out-of-order execution may be implemented (e.g., using a reorder buffer(s) and a retirement register file(s); using a future file(s), a history buffer(s), and a retirement register file(s); using a register maps and a pool of registers; etc.). The retirement unit 1054 and the physical register file(s) unit(s) 1058 are coupled to the execution cluster(s) 1060. The execution cluster(s) 1060 includes a set of one or more execution units 1062 and a set of one or more memory access units 1064. The execution units 1062 may perform various operations (e.g., shifts, addition, subtraction, multiplication) and on various types of data (e.g., scalar floating point, packed integer, packed floating point, vector integer, vector floating point). While some embodiments may include a number of execution units dedicated to specific functions or sets of functions, other embodiments may include only one execution unit or multiple execution units that all perform all functions. The scheduler unit(s) 1056, physical register file(s) unit(s) 1058, and execution cluster(s) 1060 are shown as being possibly plural because certain embodiments create separate pipelines for certain types of data/operations (e.g., a scalar integer pipeline, a scalar floating point/packed integer/packed floating point/vector integer/vector floating point pipeline, and/or a memory access pipeline that each have their own scheduler unit, physical register file(s) unit, and/or execution cluster—and in the case of a separate memory access pipeline, certain embodiments are implemented in which only the execution cluster of this pipeline has the memory access unit(s) 1064). It should also be understood that where separate pipelines are used, one or more of these pipelines may be out-of-order issue/execution and the rest in-order.

The set of memory access units 1064 is coupled to the memory unit 1070, which includes a data TLB unit 1072 coupled to a data cache unit 1074 coupled to a level 2 (L2) cache unit 1076. In one exemplary embodiment, the memory access units 1064 may include a load unit, a store address unit, and a store data unit, each of which is coupled to the data TLB unit 1072 in the memory unit 1070. The instruction cache unit 1034 is further coupled to a level 2 (L2) cache unit 1076 in the memory unit 1070. The L2 cache unit 1076 is coupled to one or more other levels of cache and eventually to a main memory.

By way of example, the exemplary register renaming, out-of-order issue/execution core architecture may implement the pipeline 1000 as follows: 1) the instruction fetch unit 1038 performs the fetch and length decoding stages 1002 and 1004; 2) the decode unit 1040 performs the decode stage 1006; 3) the rename/allocator unit 1052 performs the allocation stage 1008 and renaming stage 1010; 4) the scheduler unit(s) 1056 performs the schedule stage 1012; 5) the physical register file(s) unit(s) 1058 and the memory unit 1070 perform the register read/memory read stage 1014; the execution cluster 1060 perform the execute stage 1016; 6) the memory unit 1070 and the physical register file(s) unit(s) 1058 perform the write back/memory write stage 1018; 7) various units may be involved in the exception handling stage 1022; and 8) the retirement unit 1054 and the physical register file(s) unit(s) 1058 perform the commit stage 1024.

The core 1090 may support one or more instructions sets (e.g., the x86 instruction set (with some extensions that have

been added with newer versions); the MIPS instruction set of MIPS Technologies of Sunnyvale, Calif.; the ARM instruction set (with optional additional extensions such as NEON) of ARM Holdings of Sunnyvale, Calif.), including the instruction(s) described herein. In one embodiment, the core **1090** includes logic to support a packed data instruction set extension (e.g., AVX1, AVX2, and/or some form of the generic vector friendly instruction format (U=0 and/or U=1) previously described), thereby allowing the operations used by many multimedia applications to be performed using packed data.

It should be understood that the core may support multithreading (executing two or more parallel sets of operations or threads), and may do so in a variety of ways including time sliced multithreading, simultaneous multithreading (where a single physical core provides a logical core for each of the threads that physical core is simultaneously multithreading), or a combination thereof (e.g., time sliced fetching and decoding and simultaneous multithreading thereafter such as in the Intel® Hyperthreading technology).

While register renaming is described in the context of out-of-order execution, it should be understood that register renaming may be used in an in-order architecture. While the illustrated embodiment of the processor also includes separate instruction and data cache units **1034/1074** and a shared L2 cache unit **1076**, alternative embodiments may have a single internal cache for both instructions and data, such as, for example, a Level 1 (L1) internal cache, or multiple levels of internal cache. In some embodiments, the system may include a combination of an internal cache and an external cache that is external to the core and/or the processor. Alternatively, all of the cache may be external to the core and/or the processor.

Specific Exemplary In-Order Core Architecture

FIGS. **11A-B** illustrate a block diagram of a more specific exemplary in-order core architecture, which core would be one of several logic blocks (including other cores of the same type and/or different types) in a chip. The logic blocks communicate through a high-bandwidth interconnect network (e.g., a ring network) with some fixed function logic, memory I/O interfaces, and other necessary I/O logic, depending on the application.

FIG. **11A** is a block diagram of a single processor core, along with its connection to the on-die interconnect network **1102** and with its local subset of the Level 2 (L2) cache **1104**, according to embodiments of the invention. In one embodiment, an instruction decoder **1100** supports the x86 instruction set with a packed data instruction set extension. An L1 cache **1106** allows low-latency accesses to cache memory into the scalar and vector units. While in one embodiment (to simplify the design), a scalar unit **1108** and a vector unit **1110** use separate register sets (respectively, scalar registers **1112** and vector registers **1114**) and data transferred between them is written to memory and then read back in from a level 1 (L1) cache **1106**, alternative embodiments of the invention may use a different approach (e.g., use a single register set or include a communication path that allow data to be transferred between the two register files without being written and read back).

The local subset of the L2 cache **1104** is part of a global L2 cache that is divided into separate local subsets, one per processor core. Each processor core has a direct access path to its own local subset of the L2 cache **1104**. Data read by a processor core is stored in its L2 cache subset **1104** and can be accessed quickly, in parallel with other processor cores accessing their own local L2 cache subsets. Data written by a processor core is stored in its own L2 cache subset **1104**

and is flushed from other subsets, if necessary. The ring network ensures coherency for shared data. The ring network is bi-directional to allow agents such as processor cores, L2 caches and other logic blocks to communicate with each other within the chip. Each ring data-path is 1012-bits wide per direction.

FIG. **11B** is an expanded view of part of the processor core in FIG. **11A** according to embodiments of the invention. FIG. **11B** includes an L1 data cache **1106A** part of the L1 cache **1104**, as well as more detail regarding the vector unit **1110** and the vector registers **1114**. Specifically, the vector unit **1110** is a 16-wide vector processing unit (VPU) (see the 16-wide ALU **1128**), which executes one or more of integer, single-precision float, and double-precision float instructions. The VPU supports swizzling the register inputs with swizzle unit **1120**, numeric conversion with numeric convert units **1122A-B**, and replication with replication unit **1124** on the memory input. Write mask registers **1126** allow predicated resulting vector writes.

Processor with Integrated Memory Controller and Graphics

FIG. **12** is a block diagram of a processor **1200** that may have more than one core, may have an integrated memory controller, and may have integrated graphics according to embodiments of the invention. The solid lined boxes in FIG. **12** illustrate a processor **1200** with a single core **1202A**, a system agent **1210**, a set of one or more bus controller units **1216**, while the optional addition of the dashed lined boxes illustrates an alternative processor **1200** with multiple cores **1202A-N**, a set of one or more integrated memory controller unit(s) **1214** in the system agent unit **1210**, and special purpose logic **1208**.

Thus, different implementations of the processor **1200** may include: 1) a CPU with the special purpose logic **1208** being integrated graphics and/or scientific (throughput) logic (which may include one or more cores), and the cores **1202A-N** being one or more general purpose cores (e.g., general purpose in-order cores, general purpose out-of-order cores, a combination of the two); 2) a coprocessor with the cores **1202A-N** being a large number of special purpose cores intended primarily for graphics and/or scientific (throughput); and 3) a coprocessor with the cores **1202A-N** being a large number of general purpose in-order cores. Thus, the processor **1200** may be a general-purpose processor, coprocessor or special-purpose processor, such as, for example, a network or communication processor, compression engine, graphics processor, GPGPU (general purpose graphics processing unit), a high-throughput many integrated core (MIC) coprocessor (including 30 or more cores), embedded processor, or the like. The processor may be implemented on one or more chips. The processor **1200** may be a part of and/or may be implemented on one or more substrates using any of a number of process technologies, such as, for example, BiCMOS, CMOS, or NMOS.

The memory hierarchy includes one or more levels of cache within the cores, a set or one or more shared cache units **1206**, and external memory (not shown) coupled to the set of integrated memory controller units **1214**. The set of shared cache units **1206** may include one or more mid-level caches, such as level 2 (L2), level 3 (L3), level 4 (L4), or other levels of cache, a last level cache (LLC), and/or combinations thereof. While in one embodiment a ring based interconnect unit **1212** interconnects the integrated graphics logic **1208**, the set of shared cache units **1206**, and the system agent unit **1210**/integrated memory controller unit(s) **1214**, alternative embodiments may use any number of well-known techniques for interconnecting such units. In

one embodiment, coherency is maintained between one or more cache units **1206** and cores **1202-A-N**.

In some embodiments, one or more of the cores **1202A-N** are capable of multi-threading. The system agent **1210** includes those components coordinating and operating cores **1202A-N**. The system agent unit **1210** may include for example a power control unit (PCU) and a display unit. The PCU may be or include logic and components needed for regulating the power state of the cores **1202A-N** and the integrated graphics logic **1208**. The display unit is for driving one or more externally connected displays.

The cores **1202A-N** may be homogenous or heterogeneous in terms of architecture instruction set; that is, two or more of the cores **1202A-N** may be capable of execution the same instruction set, while others may be capable of executing only a subset of that instruction set or a different instruction set.

Exemplary Computer Architectures

FIGS. **13-16** are block diagrams of exemplary computer architectures. Other system designs and configurations known in the arts for laptops, desktops, handheld PCs, personal digital assistants, engineering workstations, servers, network devices, network hubs, switches, embedded processors, digital signal processors (DSPs), graphics devices, video game devices, set-top boxes, micro controllers, cell phones, portable media players, hand held devices, and various other electronic devices, are also suitable. In general, a huge variety of systems or electronic devices capable of incorporating a processor and/or other execution logic as disclosed herein are generally suitable.

Referring now to FIG. **13**, shown is a block diagram of a system **1300** in accordance with one embodiment of the present invention. The system **1300** may include one or more processors **1310**, **1315**, which are coupled to a controller hub **1320**. In one embodiment the controller hub **1320** includes a graphics memory controller hub (GMCH) **1390** and an Input/Output Hub (IOH) **1350** (which may be on separate chips); the GMCH **1390** includes memory and graphics controllers to which are coupled memory **1340** and a coprocessor **1345**; the IOH **1350** is couples input/output (I/O) devices **1360** to the GMCH **1390**. Alternatively, one or both of the memory and graphics controllers are integrated within the processor (as described herein), the memory **1340** and the coprocessor **1345** are coupled directly to the processor **1310**, and the controller hub **1320** in a single chip with the IOH **1350**.

The optional nature of additional processors **1315** is denoted in FIG. **13** with broken lines. Each processor **1310**, **1315** may include one or more of the processing cores described herein and may be some version of the processor **1200**.

The memory **1340** may be, for example, dynamic random access memory (DRAM), phase change memory (PCM), or a combination of the two. For at least one embodiment, the controller hub **1320** communicates with the processor(s) **1310**, **1315** via a multi-drop bus, such as a frontside bus (FSB), point-to-point interface such as QuickPath Interconnect (QPI), or similar connection **1395**.

In one embodiment, the coprocessor **1345** is a special-purpose processor, such as, for example, a high-throughput MIC processor, a network or communication processor, compression engine, graphics processor, GPGPU, embedded processor, or the like. In one embodiment, controller hub **1320** may include an integrated graphics accelerator.

There can be a variety of differences between the physical resources **1310**, **1315** in terms of a spectrum of metrics of

merit including architectural, microarchitectural, thermal, power consumption characteristics, and the like.

In one embodiment, the processor **1310** executes instructions that control data processing operations of a general type. Embedded within the instructions may be coprocessor instructions. The processor **1310** recognizes these coprocessor instructions as being of a type that should be executed by the attached coprocessor **1345**. Accordingly, the processor **1310** issues these coprocessor instructions (or control signals representing coprocessor instructions) on a coprocessor bus or other interconnect, to coprocessor **1345**. Coprocessor (s) **1345** accept and execute the received coprocessor instructions.

Referring now to FIG. **14**, shown is a block diagram of a first more specific exemplary system **1400** in accordance with an embodiment of the present invention. As shown in FIG. **14**, multiprocessor system **1400** is a point-to-point interconnect system, and includes a first processor **1470** and a second processor **1480** coupled via a point-to-point interconnect **1450**. Each of processors **1470** and **1480** may be some version of the processor **1200**. In one embodiment of the invention, processors **1470** and **1480** are respectively processors **1310** and **1315**, while coprocessor **1438** is coprocessor **1345**. In another embodiment, processors **1470** and **1480** are respectively processor **1310** coprocessor **1345**.

Processors **1470** and **1480** are shown including integrated memory controller (IMC) units **1472** and **1482**, respectively. Processor **1470** also includes as part of its bus controller units point-to-point (P-P) interfaces **1476** and **1478**; similarly, second processor **1480** includes P-P interfaces **1486** and **1488**. Processors **1470**, **1480** may exchange information via a point-to-point (P-P) interface **1450** using P-P interface circuits **1478**, **1488**. As shown in FIG. **14**, IMCs **1472** and **1482** couple the processors to respective memories, namely a memory **1432** and a memory **1434**, which may be portions of main memory locally attached to the respective processors.

Processors **1470**, **1480** may each exchange information with a chipset **1490** via individual P-P interfaces **1452**, **1454** using point to point interface circuits **1476**, **1494**, **1486**, **1498**. Chipset **1490** may optionally exchange information with the coprocessor **1438** via a high-performance interface **1439**. In one embodiment, the coprocessor **1438** is a special-purpose processor, such as, for example, a high-throughput MIC processor, a network or communication processor, compression engine, graphics processor, GPGPU, embedded processor, or the like.

A shared cache (not shown) may be included in either processor or outside of both processors, yet connected with the processors via P-P interconnect, such that either or both processors' local cache information may be stored in the shared cache if a processor is placed into a low power mode.

Chipset **1490** may be coupled to a first bus **1416** via an interface **1496**. In one embodiment, first bus **1416** may be a Peripheral Component Interconnect (PCI) bus, or a bus such as a PCI Express bus or another third generation I/O interconnect bus, although the scope of the present invention is not so limited.

As shown in FIG. **14**, various I/O devices **1414** may be coupled to first bus **1416**, along with a bus bridge **1418** which couples first bus **1416** to a second bus **1420**. In one embodiment, one or more additional processor(s) **1415**, such as coprocessors, high-throughput MIC processors, GPGPU's, accelerators (such as, e.g., graphics accelerators or digital signal processing (DSP) units), field programmable gate arrays, or any other processor, are coupled to first bus **1416**. In one embodiment, second bus **1420** may be a low

pin count (LPC) bus. Various devices may be coupled to a second bus **1420** including, for example, a keyboard and/or mouse **1422**, communication devices **1427** and a storage unit **1428** such as a disk drive or other mass storage device which may include instructions/code and data **1430**, in one embodiment. Further, an audio I/O **1424** may be coupled to the second bus **1420**. Note that other architectures are possible. For example, instead of the point-to-point architecture of FIG. **14**, a system may implement a multi-drop bus or other such architecture.

Referring now to FIG. **15**, shown is a block diagram of a second more specific exemplary system **1500** in accordance with an embodiment of the present invention. Like elements in FIGS. **14** and **15** bear like reference numerals, and certain aspects of FIG. **14** have been omitted from FIG. **15** in order to avoid obscuring other aspects of FIG. **15**.

FIG. **15** illustrates that the processors **1470**, **1480** may include integrated memory and I/O control logic ("CL") **1472** and **1482**, respectively. Thus, the CL **1472**, **1482** include integrated memory controller units and include I/O control logic. FIG. **15** illustrates that not only are the memories **1432**, **1434** coupled to the CL **1472**, **1482**, but also that I/O devices **1514** are also coupled to the control logic **1472**, **1482**. Legacy I/O devices **1515** are coupled to the chipset **1490**.

Referring now to FIG. **16**, shown is a block diagram of a SoC **1600** in accordance with an embodiment of the present invention. Similar elements in FIG. **12** bear like reference numerals. Also, dashed lined boxes are optional features on more advanced SoCs. In FIG. **16**, an interconnect unit(s) **1602** is coupled to: an application processor **1610** which includes a set of one or more cores **202A-N** and shared cache unit(s) **1206**; a system agent unit **1210**; a bus controller unit(s) **1216**; an integrated memory controller unit(s) **1214**; a set of one or more coprocessors **1620** which may include integrated graphics logic, an image processor, an audio processor, and a video processor; an static random access memory (SRAM) unit **1630**; a direct memory access (DMA) unit **1632**; and a display unit **1640** for coupling to one or more external displays. In one embodiment, the coprocessor(s) **1620** include a special-purpose processor, such as, for example, a network or communication processor, compression engine, GPGPU, a high-throughput MIC processor, embedded processor, or the like.

Embodiments of the mechanisms disclosed herein may be implemented in hardware, software, firmware, or a combination of such implementation approaches. Embodiments of the invention may be implemented as computer programs or program code executing on programmable systems comprising at least one processor, a storage system (including volatile and non-volatile memory and/or storage elements), at least one input device, and at least one output device.

Program code, such as code **1430** illustrated in FIG. **14**, may be applied to input instructions to perform the functions described herein and generate output information. The output information may be applied to one or more output devices, in known fashion. For purposes of this application, a processing system includes any system that has a processor, such as, for example; a digital signal processor (DSP), a microcontroller, an application specific integrated circuit (ASIC), or a microprocessor.

The program code may be implemented in a high level procedural or object oriented programming language to communicate with a processing system. The program code may also be implemented in assembly or machine language, if desired. In fact, the mechanisms described herein are not

limited in scope to any particular programming language. In any case, the language may be a compiled or interpreted language.

One or more aspects of at least one embodiment may be implemented by representative instructions stored on a machine-readable medium which represents various logic within the processor, which when read by a machine causes the machine to fabricate logic to perform the techniques described herein. Such representations, known as "IP cores" may be stored on a tangible, machine readable medium and supplied to various customers or manufacturing facilities to load into the fabrication machines that actually make the logic or processor.

Such machine-readable storage media may include, without limitation, non-transitory, tangible arrangements of articles manufactured or formed by a machine or device, including storage media such as hard disks, any other type of disk including floppy disks, optical disks, compact disk read-only memories (CD-ROMs), compact disk rewritable's (CD-RWs), and magneto-optical disks, semiconductor devices such as read-only memories (ROMs), random access memories (RAMs) such as dynamic random access memories (DRAMs), static random access memories (SRAMs), erasable programmable read-only memories (EPROMs), flash memories, electrically erasable programmable read-only memories (EEPROMs), phase change memory (PCM), magnetic or optical cards, or any other type of media suitable for storing electronic instructions.

Accordingly, embodiments of the invention also include non-transitory, tangible machine-readable media containing instructions or containing design data, such as Hardware Description Language (HDL), which defines structures, circuits, apparatuses, processors and/or system features described herein. Such embodiments may also be referred to as program products.

Emulation (Including Binary Translation, Code Morphing, Etc.)

In some cases, an instruction converter may be used to convert an instruction from a source instruction set to a target instruction set. For example, the instruction converter may translate (e.g., using static binary translation, dynamic binary translation including dynamic compilation), morph, emulate, or otherwise convert an instruction to one or more other instructions to be processed by the core. The instruction converter may be implemented in software, hardware, firmware, or a combination thereof. The instruction converter may be on processor, off processor, or part on and part off processor.

FIG. **17** is a block diagram contrasting the use of a software instruction converter to convert binary instructions in a source instruction set to binary instructions in a target instruction set according to embodiments of the invention. In the illustrated embodiment, the instruction converter is a software instruction converter, although alternatively the instruction converter may be implemented in software, firmware, hardware, or various combinations thereof. FIG. **17** shows a program in a high level language **1702** may be compiled using an x86 compiler **1704** to generate x86 binary code **1706** that may be natively executed by a processor with at least one x86 instruction set core **1716**. The processor with at least one x86 instruction set core **1716** represents any processor that can perform substantially the same functions as an Intel processor with at least one x86 instruction set core by compatibly executing or otherwise processing (1) a substantial portion of the instruction set of the Intel x86 instruction set core or (2) object code versions of applications or other software targeted to run on an Intel processor

29

with at least one x86 instruction set core, in order to achieve substantially the same result as an Intel processor with at least one x86 instruction set core. The x86 compiler **1704** represents a compiler that is operable to generate x86 binary code **1706** (e.g., object code) that can, with or without additional linkage processing, be executed on the processor with at least one x86 instruction set core **1716**. Similarly, FIG. **17** shows the program in the high level language **1702** may be compiled using an alternative instruction set compiler **1708** to generate alternative instruction set binary code **1710** that may be natively executed by a processor without at least one x86 instruction set core **1714** (e.g., a processor with cores that execute the MIPS instruction set of MIPS Technologies of Sunnyvale, Calif. and/or that execute the ARM instruction set of ARM Holdings of Sunnyvale, Calif.). The instruction converter **1712** is used to convert the x86 binary code **1706** into code that may be natively executed by the processor without an x86 instruction set core **1714**. This converted code is not likely to be the same as the alternative instruction set binary code **1710** because an instruction converter capable of this is difficult to make; however, the converted code will accomplish the general operation and be made up of instructions from the alternative instruction set. Thus, the instruction converter **1712** represents software, firmware, hardware, or a combination thereof that, through emulation, simulation or any other process, allows a processor or other electronic device that does not have an x86 instruction set processor or core to execute the x86 binary code **1706**.

What is claimed is:

1. A method comprising:

executing a single instruction to, for each data lane of a source vector register, calculate a horizontal and cross addition or subtraction between packed data elements of the source vector register where the decision of whether to add or subtract is dependent upon a bit in an immediate value of the instruction, wherein the instruction includes a destination vector register operand, the source vector register operand, the immediate, and an opcode; and

storing a result of each a horizontal and cross addition or subtraction between packed data elements into the destination register, wherein for the least significant data element position of the destination register, the result stored is the least significant data element of the source register being added to or subtracted from the second most least significant data element of the source register wherein the decision of addition or subtraction is based on the least significant bit of the immediate;

wherein for the second most least significant data element position of the destination register, the result stored is the result of the third most least significant data element of the source register being added to or subtracted from the fourth most least significant data element of the source register wherein the decision of addition or subtraction is based on the third most least significant bit of the immediate;

wherein for the third most least significant data element position of the destination register, the result stored is the result of the second most least significant data element of the source register being added to or subtracted from the least significant data element of the source register wherein the decision of addition or subtraction is based on the second most least significant bit of the immediate; and

wherein for the fourth most significant data element position of the destination register, the result stored is

30

the result of the fourth most least significant data element of the source register being added to or subtracted from the third most least significant data element of the source register wherein the decision of addition or subtraction is based on the fourth most least significant bit of the immediate.

2. The method of claim 1, wherein there are a plurality of data lanes.

3. The method of claim 1, wherein a number of data lanes to be processed is dependent upon the size of the destination vector register.

4. The method of claim 1, wherein the source and destination vector registers are 128-bit, 256-bit, or 512-bit in size.

5. The method of claim 1, wherein the packed data elements of the source and destination registers are 8-bit, 16-bit, 32-bit, or 64-bit in size.

6. The method of claim 5, wherein the size of the packed data elements of the source is defined by the opcode.

7. The method of claim 1, wherein the immediate is an 8-bit value.

8. An article of manufacture comprising:

a non-transitory machine-readable storage medium having stored thereon an occurrence of a single instruction, wherein the instruction's format specifies as its source operands a vector register and an immediate and specifies as its destination a single destination vector register, and wherein the instruction format includes an opcode which instructs a machine, responsive to the single occurrence of the single instruction, to cause for each data lane of the source vector register, a calculation of a horizontal and cross addition or subtraction between packed data elements of the source vector register where the decision of whether to add or subtract is dependent upon a bit in the immediate value of the instruction, and storage of a result of each a horizontal and cross addition or subtraction between packed data elements results into the destination register,

wherein for the least significant data element position of the destination register, the result stored is the least significant data element of the source register being added to or subtracted from the second most least significant data element of the source register wherein the decision of addition or subtraction is based on the least significant bit of the immediate;

wherein for the second most least significant data element position of the destination register, the result stored is the result of the third most least significant data element of the source register being added to or subtracted from the fourth most least significant data element of the source register wherein the decision of addition or subtraction is based on the third most least significant bit of the immediate;

wherein for the third most least significant data element position of the destination register, the result stored is the result of the second most least significant data element of the source register being added to or subtracted from the least significant data element of the source register wherein the decision of addition or subtraction is based on the second most least significant bit of the immediate; and

wherein for the fourth most significant data element position of the destination register, the result stored is the result of the fourth most least significant data element of the source register being added to or subtracted from the third most least significant data element of the source register wherein the decision of

31

addition or subtraction is based on the fourth most least significant bit of the immediate.

9. The article of manufacture of claim 8, wherein there are a plurality of data lanes.

10. The article of manufacture of claim 8, wherein a number of data lanes to be processed is dependent upon the size of the destination vector register.

11. The article of manufacture of claim 8, wherein the source and destination vector registers are 128-bit, 256-bit, or 512-bit in size.

12. The article of manufacture of claim 8, wherein the packed data elements of the source and destination registers are 8-bit, 16-bit, 32-bit, or 64-bit in size.

13. The article of manufacture of claim 8, wherein the size of the packed data elements of the source is defined by the opcode.

14. The article of manufacture of claim 8, wherein the immediate is an 8-bit value.

15. An apparatus comprising;

a hardware decoder to decode a single instruction that includes a destination vector register operand, a source vector register operand, an immediate, and an opcode; execution circuitry to, for each data lane of the source vector register, calculate a horizontal and cross addition or subtraction between packed data elements of the source vector register where the decision of whether to add or subtract is dependent upon a bit in the immediate value of the instruction, and store a result for each a horizontal and cross addition or subtraction between packed data elements results into the destination register, wherein for the least significant data element position of the destination register, the result stored is the least significant data element of the source register being added to or subtracted from the second most least significant data element of the source register wherein the decision of addition or subtraction is based on the least significant bit of the immediate;

32

wherein for the second most least significant data element position of the destination register, the result stored is the result of the third most least significant data element of the source register being added to or subtracted from the fourth most least significant data element of the source register wherein the decision of addition or subtraction is based on the third most least significant bit of the immediate;

wherein for the third most least significant data element position of the destination register, the result stored is the result of the second most least significant data element of the source register being added to or subtracted from the least significant data element of the source register wherein the decision of addition or subtraction is based on the second most least significant bit of the immediate; and

wherein for the fourth most significant data element position of the destination register, the result stored is the result of the fourth most least significant data element of the source register being added to or subtracted from the third most least significant data element of the source register wherein the decision of addition or subtraction is based on the fourth most least significant bit of the immediate.

16. The apparatus of claim 15, wherein the source and destination vector registers are 128-bit, 256-bit, or 512-bit in size.

17. The apparatus of claim 15, wherein the packed data elements of the source and destination registers are 8-bit, 16-bit, 32-bit, or 64-bit in size.

18. The apparatus of claim 15, wherein the size of the packed data elements of the source is defined by the opcode.

19. The apparatus of claim 15, wherein the immediate is an 8-bit value.

* * * * *